

# The Great Distractor

James Bridle



# About the Author

## James Bridle

Writer and Artist

James Bridle is a writer and artist based in London. Their writing on literature, culture and networks has been featured in publications including Wired, the Atlantic, the New Statesman, the Guardian, and the Observer. They are the author of 'New Dark Age: Technology and the End of the Future' (2018) and 'Ways of Being: Beyond Human Intelligence' (2022), and they wrote and presented "New Ways of Seeing" for BBC Radio 4 in 2019. Bridle holds a Master's Degree in Computer Science and Cognitive Science from University College, London.

# The Great Distractor

The children's television show *Teletubbies* premiered on British TV on 31st March 1997. It immediately attracted controversy. Initially, this controversy stemmed from the displacement of another much-loved children's show from the airwaves, but as the show gained in popularity, the media started paying attention. By August it had a regular audience of two million – a viewership that was unheard of for a programme aimed at toddlers.

The titular Teletubbies were a group of four multi-coloured creatures with televisions in their abdomens and antennae protruding from their heads: Tinky Winky, Dipsy, Laa-Laa, and Po. They lived in a hobbit-like dome in a green valley, otherwise populated only by rabbits, conversational trumpets, a sentient vacuum cleaner, and a human baby who lived in the sun. Twice per episode, the Teletubbies would watch a segment about real-world toddlers on their belly screens, in between meals of Tubby Toast and jaunts around the verdant countryside. So far, so pastoral.

What set *Teletubbies* apart from other much-loved children's programming was the behaviour of its protagonists: they spoke in a pre-verbal “goo goo” style, engaged in apparently meaningless rituals, and had a pronounced tendency to repeat themselves, as well as whole sections of the programme. This affronted the parents of many young children (although the actual number of complaints was likely much exaggerated by the tabloid press) who were concerned that such nonsense would inhibit their children's own development: “Teletubbies is boring and repetitive,” fumed one parent. “The characters have ridiculous names and don't speak coherently.”<sup>2</sup>



Fig. 1<sup>1</sup>

In fact, one of the show's co-creators, Ian Davenport, had a background in speech therapy, and his partner Anne Wood had a long history and a string of hits in children's TV. Their intention was to create a world that was both warm and welcoming to two-to-four year olds, and which prepared them for interactions with high technology: hence the televisions. The characters were deliberately unthreatening as well as slightly slower than actual toddlers, who in turn felt empathy and a desire to join in and help. Likewise, the constant repetition which infuriated adults instead delighted children, who felt security and confidence in their ability to predict what was going to happen next: a critical component of early learning.<sup>3</sup> The surreal, even psychedelic aspects of the programme also helped build a dedicated following among students and other stoners – a significant proportion of that two million-strong audience, and another stick used by the media to bash it.)

There is an implicit expectation that children's programming, if it is going to exert such a strong control over the minds of young children, ought at least to be educational

But behind the complaints about nonsense language and irritating music, what most disturbed parents and others about *Teletubbies* was just how damn good it was at doing what it did: capturing and holding the attention of young, often pre-verbal children. Placed in the glow of its primary-coloured Eden they would become utterly transfixed, as to no show before. Almost immediately they would fall into sync with its rhythms, pick up its catchphrases, join in its dances, and most disconcertingly of all, demand in unison with the four protagonists at each show's end: "Again! Again! Again!"

There is an implicit expectation that children's programming, if it is going to exert such a strong control over the minds of young children, ought at least to be educational. Although hardly alone in this, *Teletubbies* seemed to violate this principle, despite its spellbinding command. This expectation, which did not arise fully formed at the birth of television, is largely due to *Sesame Street*, the groundbreaking U.S.-based series and cultural phenomenon which debuted in 1969 and likewise achieved mass popularity in a few short months. (In a 1970 *New Yorker* cartoon, a policeman approaches a mother and daughter in the park to ask, "Why isn't that child at home watching 'Sesame Street'?"<sup>4</sup>)

While Anne Wood got the initial idea for *Teletubbies* from a dream, *Sesame Street* emerged from a moment of clarity and a long process of research and development. One morning in 1965, the Vice-President of the Carnegie Corporation Lloyd Morisett entered his living room to find his three-year-old daughter sprawled on the floor in her pyjamas, engrossed in the test pattern. It was another half hour before the networks



Fig. 2<sup>5</sup>

started broadcasting, yet across the country, thousands, even millions of children were looking for something to watch. Morisett realised there was an opportunity to do something useful with all that screen time, and commissioned a report from TV Producer Joan Ganz Cooney: “The Potential Uses of Television in Preschool Education”. Cooney noted that “nearly half the nation’s school districts do not now have kindergartens,” but “more households have televisions than bathtubs, telephones, vacuum cleaners, toasters, or a regular daily newspaper.” Moreover, researchers universally agreed on the power of television to attract young minds, while lamenting the state of its programming: “Anyone who has small television viewers at home can testify to the fascination that commercials hold for children. Parents report that their children learn to recite all sorts of advertising slogans, read product names on the screen (and, more remarkably, elsewhere), and to sing commercial jingles. It is of course open to serious question how valuable the content is that these commercials teach, but they do prove a point; children can and do learn, in the traditional educational sense, from watching television.”<sup>6</sup>

Following Cooney’s report, the Children’s Television Workshop was set up to develop and make what eventually became *Sesame Street*. Two psychologists from Harvard, Edward L. Palmer and Gerald S. Lesser, were appointed to expand on the report’s research. *Sesame Street* was the first educational programme to have a curriculum, formulated in day-long seminars with child development experts, and the first to fund extensive research into its effectiveness and outcomes.

One of the workshop’s key innovations was a device Palmer termed “the distractor”. Initially on visits to “poor children in day-care centres” and later in



specially designed observation laboratories, Palmer and his colleagues would set up a television showing *Sesame Street* episodes, and next to it they would set up a slide projector, which would show a new colour image every seven and a half seconds.<sup>7</sup> “We had the most varied set of slides we could imagine,” said Palmer. “We would have a body riding down the street with his arms out, a picture of a tall building, a leaf floating through ripples of water, a rainbow, a picture taken through a microscope, an Escher drawing. Anything to be novel, that was the idea.”<sup>8</sup> Observers would sit at the side of the room and watch a group of children as the episodes played, noting what held their attention and what didn’t every six and a half seconds for up to an hour at a time.

Individual characters, short skits, and even whole episodes lived or died by the distractor. By the time they got to broadcast, the show’s attention rating averaged around 85 to 90 percent, and some reached 100 percent, as the crew learned to trust the distractor and absorbed its lessons. For Palmer, the distractor enabled the team to hone *Sesame Street* into a finely tuned educational instrument: “After the third or fourth season, I’d say it was rare that we ever had a segment below eighty five percent. We would almost never see something in the fifty to sixty percent range, and if we did, we’d fix it. You know Darwin’s terms about the survival of the fittest? We had a mechanism to identify the fittest and decide what should survive.”

*Sesame Street* was pitched at a slightly higher age bracket than *Teletubbies*; the Muppets were and remain a mostly smart, funny, and prolifically verbal gang, while the *Teletubbies* are pre-literate and buffoonish. (Big Bird, a character modelled on an easily flustered four-year-old who needs a lot of help, is the Muppet closest aesthetically and emotionally to the *Teletubbies*.) And unlike *Sesame Street*, which combined extensive surveillance and data-driven feedback with an earnest, patrician pedagogical agenda, *Teletubbies* was not based on rigorous developmental theory or any higher ambition than providing stimulating, good-natured entertainment.



Fig. 3<sup>9</sup>

Nevertheless, some aspects of *Teletubbies* still seem progressive. While *Sesame Street* included Black and Latinx characters from its nineteen-seventies beginnings, British children's television in the nineties was still much less diverse. The Teletubbies were multiracial, with the producers stating that the darker-skinned Dipsy was Black ("a furry, funky Jamaican toddler"<sup>10</sup>), and the diminutive Po was Chinese, as were the respective – and invisible – actors who played them. Tinky Winky proved to be the catalyst for the show's fiercest controversy, when the American televangelist Jerry Falwell claimed he was gay. "He is purple – the gay-pride color; and his antenna is shaped like a triangle – the gay-pride symbol," he wrote. He also noted that Tinky Winky carries a handbag, a sure sign of sexual deviancy. The BBC's official response stated that "Tinky Winky is simply a sweet, technological baby with a magic bag"<sup>11</sup> but the idea stuck in the public consciousness. That a majority approved of or didn't care about Tinky-Winky's orientation, however notional, was a sign of changing times.

Other memories of the programme seem a little darker. The giant rabbits, imported from Holland, had enlarged hearts and kept dying on set. The show's extraordinary popularity led to a siege-like mentality. At one point, the producers had to hire security guards and build a special changing tent to prevent photos of half-dressed, decapitated Teletubbies with the actors' heads visible from leaking out to the media. There was talk of imposing a no-fly zone over Tubbyland.<sup>12</sup> Once fully clothed, the actors were blind and almost deaf: instructions ("turn left, walk forward three paces") had to be transmitted to earpieces. They worked in darkness and frequently overheated.<sup>13</sup>

It's that last detail which brings to mind the Mechanical Turk, that oft-deployed symbol of mixed human and machine agency. This celebrated automaton toured the courts of Europe throughout the late eighteenth and early nineteenth centuries, defeating all comers at a game of chess, apparently conducted – very effectively – by a robot dressed in a robe and fez. A Teletubby *avant la lettre*, the Mechanical Turk foregrounded a mysterious and glamorous technology, while concealing more complex, yet all too human labours within.

The Mechanical Turk was constructed and unveiled in 1770 by Hungarian inventor Wolfgang von Kempelen, who presented it to Empress Maria Theresa of Austria. For eight decades it surprised, delighted, and mystified audiences, playing and winning matches against public figures such as Napoleon Bonaparte and Benjamin Franklin. Only after its loss in a fire in 1850 was its secret revealed: a succession of the world's greatest chess players, several of them

drunk or indebted, crouched in the wooden cabinet beneath the board. A sliding seat allowed them to move from side to side when doors were opened to give the illusion that the cabinet was empty, while a small chimney pipe conducted the smoke of a single candle up and out of the Turk's turban.

The Turk remains emblematic of the only real ghost in the machine: human agency. As computational systems have become ever more complex and embedded, often invisibly, into everyday life, the work of opening the cabinet to reveal the real actors and intentions beneath the opaque and often confusing exterior becomes more and more critical. The ability to disentangle the operator from the machine will be crucial to understanding what has happened to the techniques so carefully crafted by the Children's Television Workshop and so entrancingly deployed by the Teletubbies, and the ways in which they are deployed today.

*Teletubbies* was technological – and even uncanny – from the outset, as producer Anne Wood recalled of its inception: “We were trying to reflect children's experience back to them. It was a world in which technology was developing at an extraordinary rate: children were growing up in an environment where



Fig. 4<sup>14</sup>

objects spoke from the walls or moved about.” In the revival of the series in 2015, *Teletubbies* introduced the Tubby Phone, a nod to post-television technologies. The arrival of the Tubby Phone hints at the vast shift taking place in media technologies off-screen in the late nineties and early noughties: the explosive

growth of the internet and its various subsumed devices. The Tubby Phone could be used to make calls and take photos, just like a contemporary smartphone, although its primary use was to trigger the Tubby Phone Dance, a cavort around the Tubbytronic Superdome to the refrain of “Ring, ring, ring.”

*Sesame Street* and *Teletubbies* used different methods to get the attention of their young viewers, but attention was the main thing: once gained, you could do anything with it. The advent of digital technologies democratized access to



attention, and provided amateurs, upstarts, and established companies with ways to profit from its successful acquisition and retention. And the target audience now wasn't merely Lloyd Morissett's bored pre-schoolers lounging on the living room floor on Saturday mornings, but potentially every connected human being on the planet, wherever they were and whatever they were supposed to be doing instead.

In 1997, the same year as the attention-grabbing debut of the *Teletubbies*, BJ Fogg, a psychology professor at Stanford, introduced a new topic at the SIGCHI Conference on Human Factors in Computing Systems. He called it "captology," or the study of computers as persuasive technologies.<sup>15</sup> By this, he meant the use of interactive technology to change a person's attitudes or behaviour, and he broke his new science down into a number of different areas and examples for future research. As examples of persuasive technology, he cited tools such as heart rate monitors and step counters (for changing behaviour and attitudes around fitness), media such as video games which contained educational information, and "social actors" such as Tamagotchis, the then-popular electronic pets. While the examples he found mostly operated within the domains of health and education, he foresaw that persuasive computers could be deployed to address a range of "pervasive social or personal problems," including occupational safety, environmental issues, and financial management.

The advent of digital technologies democratised access to attention, and provided amateurs, upstarts, and established companies with ways to profit from its successful

Fogg also sounded a note of caution. Captology research would require a strong ethical component, because persuasive techniques could easily be employed by those without the best interests of their users at heart. He was concerned, for example, by the possibility of casinos using captology to make their slot machines more seductive, or corporations deploying captology to track and manipulate employee behaviour. Education was the best defence against such abuses, he believed. "First, increased knowledge about persuasive computers allows people more opportunity to adopt such technologies to enhance their own lives, if they choose," he wrote. "Second, knowledge about persuasive computers helps people recognize when technologies are using tactics to persuade them."

In service of the second of these aims, Fogg published *Persuasive Technology: Using Computers to Change What We Think and Do* in 2003. He dedicated a chapter of the

book to the ethics of captology, and expanded on his concerns. Among these were succinct observations such as “the novelty of the technology can mask its persuasive intent,” “computers can be proactively persistent,” and “computers can affect emotions but can’t be affected by them.” These concerns were synthesised from patterns he saw repeated in a special class he ran at Stanford, in which students were encouraged to explore the “dark side” of persuasive technology and develop applications which violated the consent or autonomy of users.<sup>16</sup>

The intent of this class, he wrote, was to help his students “understand the implications of future technology and how to prevent unethical applications or mitigate their impact.” But the ethics class was only one part of the programme: significantly more time was spent figuring out how to apply captology in the real world.

In 2006, a couple of Fogg’s students developed a prototype application called “Send the Sunshine”. Years before smartphones became omnipresent, they imagined that one day mobile devices might be used not just to send data but



Fig. 5<sup>17</sup>

to to send emotions to other people. They envisioned a system which would prompt people enjoying good weather to share it with people in gloomier climes, in the form of a description or photograph. It’s a classic example of the bright side of captology: the use of data to subtly influence personal behaviour to make the wider world a better, happier place.

People remember Send the Sunshine because one of those students, Mike Krieger, went on to found Instagram, an application used today by more than a billion people worldwide. Instagram is a key exemplar of Fogg’s principles. Users are “persuaded” to keep using the app through the social rewards of likes and comments. At a deeper level, they’re also empowered by the more subtle choices of filters and effects users can apply to each photo. “Sure, there’s a functional benefit: the user has control over their images,” noted Fogg himself. “But the real transaction is emotional: before you even post anything, you get to feel like an artist.”<sup>18</sup>

The approach kept working. The year after the development of Send the Sunshine – which was essentially a thought experiment, albeit a prescient one

– Fogg tried a new teaching tool. Facebook, then a growing but by no means dominant force in social media, had just launched its Apps platform, a way for developers to build games, toys, and utilities for users, and Fogg saw a chance to test out some of his ideas in the real world. In 2007, working in teams of three, seventy-five students in his class were tasked with creating apps for Facebook, with the highest grades going to those whose app attracted the most users, using the principles of captology. Many of the apps were extremely simple, but the results were, and remain, extraordinary.

[P]ersuasive technologies made Facebook what it was: an intricately-tooled machine for the capture and retention of attention

Within a few weeks, one app called *KissMe* had passed a million users. *KissMe* was simple but effective: it invited users to send another user a kiss, and the kiss acted as an invitation to the app, growing its install base. Another app, *Send Hotness*, was close behind: it enabled users to rank their friends' attractiveness, and the top ten were sent invitations. Flattery, it seems, will get you everywhere.<sup>19</sup> By the end of the term, the class's apps had a total of more than sixteen million users. And more consequentially, they'd generated more than a million dollars in actual revenue, through the accompanying, also-novel medium of digital advertising.<sup>20</sup> Attention itself could be monetised, and captology provided the toolkit for its acquisition.

The “Facebook class”, as it became known, kick-started a number of successful careers. Within a few years, a significant number of the students were running multi-million, if not -billion, dollar start-ups in Silicon Valley, and the Valley took note of the approach the class embodied: small teams, iterating quickly, seeking fast money from clever hacks. Captology became key to the growth of Facebook in particular, in a process of feedback that rippled across the technology industry. Facebook's apps – and later the platform itself – enabled persuasive technologies to operate at scale, and persuasive technologies made Facebook what it was: an intricately-tooled machine for the capture and retention of attention. An entire industry began to organise itself around capturing and holding attention, and sharpening the tools for doing so.

Back in 2006, Mike Krieger's partner in Fogg's class, the other half of Send the Sunshine, was a student called Tristan Harris. Harris ended up taking a somewhat different path to his former classmate. After his own start-up, a search company called Apture, was acquired by Google, Harris found himself working inside the second-largest internet company in the world. And he didn't like what he saw.

In 2013, after a couple of years within Google, Harris circulated internally a 141-slide presentation he'd created entitled "A Call to Minimize Distraction & Respect Users' Attention".<sup>21</sup> Clearly harking back to his classes with Fogg on the opportunities and dangers of persuasive computing, Harris called out the ways in which Silicon Valley in general, and Google, Apple, and Facebook in particular, were causing actual harm by exploiting the vulnerabilities of their users. He called the companies' products "attention casinos", tapping into common human addictions like intermittent variable reward systems. This is the same psychological vulnerability that comes into play when feeding slot machines, hitting refresh on email accounts, or swiping down on news feeds. He outlined the mental and physiological effects of spending time inside stress-inducing applications, including "email apnea", a term coined by former Apple executive Linda Stone to describe the tendency to hold one's breath when reading emails. Such a reaction triggers a fight-or-flight response, raises stress levels, and degrades decision making.<sup>22</sup> Harris also showed how the drive towards making interactions and experiences friction-free and "seamless" removed the time for reflection and thoughtfulness necessary for making good decisions – or to exercise free will at all.

Harris' aim wasn't just to castigate his colleagues. He genuinely believed that they could change the way they worked – and he pointed out how they already did this internally. A photo of a snack bar within the Google headquarters itself showed that it was designed to include all kinds of food, but that the treats deemed unhealthy were placed inside jars at the top of the shelving, making them that bit harder to reach than the healthier options below. Googlers could apply the same technique to their products he thought: increasing the friction just a little bit, to give users time to breathe and think, rather than trying to trick or hack them into actions which were against their better interests. "Never before in history have the decisions of a handful of designers (mostly men, white, living in SF, aged 25–35) working at 3 companies," he wrote, "had so much impact on how millions of people around the world spend their attention ... We should feel an enormous responsibility to get this right."

Harris believed Google was in a unique position to address these issues. It was a problem which was beyond the abilities of "niche startups" to address, and could only be tackled "top-down, from large institutions that define the standards for millions of people." Google, he noted, set those standards for more than half of the world's mobile phones, and carried out "> 11 billion interruptions to people's lives *every day...* (this is nuts!)". But he also noted – without the emphasis on Google –



that “successful products compete by exploiting these vulnerabilities, so they can’t remove them without sacrificing their success or growth.” In a corporate culture – indeed, on a planet – utterly shaped by the profit motive, and entirely in thrall to shareholder value, which approach do you think was likely to win out?

In 2018, Google launched a new suite of tools on Android to help people manage their technology use better, and a new website, [wellbeing.google](https://wellbeing.google), with tips for focussing one’s time, and unplugging more often. Many news outlets reported that this new direction had been prompted by Harris’ report back in 2013.<sup>23</sup> If so, they weren’t acknowledging it. Although he was named Google’s first “design ethicist” after the report reached the top levels of the company, Harris left the company in 2016 having failed to convince managers to actually change their products.

To be clear:  
YouTube decides  
what most people  
watch.

Moreover, Google’s tweaks left most of their most powerful techniques untouched. In his original presentation, Harris gave a number of examples of irresponsible design, subtly altering screenshots of Google and other companies’ products to show what was really happening. One example was the notification which Facebook users receive when they’re tagged in a photo: an email telling them they’ve been tagged, and two buttons with the words “See photo” and “Go to notifications”. He suggested the first one should be rewritten as “Spend next 20 minutes on FB” to acknowledge that the designers weren’t really doing the user a useful service, but were attempting to hijack their attention. Another example was YouTube’s “Up Next” sidebar of recommended videos. Harris replaced the titles of these with the real effects of being caught in the endless slipstream of suggestions: “Don’t sleep til 2am”, “Don’t sleep til 3am”, “Feel sick tomorrow morning”, “Wake up feeling exhausted”, and “Regret staying up all night watching YouTube”.

YouTube’s recommendation system counts as one of the most egregious examples of what Harris has called “the race to the bottom of the brain stem”: the efforts by Silicon Valley companies to hijack human attention at the source. Over the years, this crucial component of Google’s video-on-demand service has morphed from a simple keyword associator into a sophisticated, AI-driven analyser, classifier, connector, and even commissioner of content.





How much power does this system have, and what does it mean to have that power? YouTube has two billion active users every month, more than any other social media platform or video-on-demand platform.<sup>24</sup> Collectively, those users watch billions upon billions of videos. And of those videos, 70 percent are selected for them by YouTube's recommendation engine. What is watched on YouTube is overwhelmingly driven by Google's finely-honed persuasion system.<sup>25</sup> To be clear: YouTube *decides* what most people watch. Moreover, because YouTube knows what people watch, it can seed that information to those who make videos, creating a feedback loop between demand and supply, where the algorithm doesn't merely select content for people to watch, but actively participates in the generation of that content, ensuring that there is a steady stream of new recommendations to keep people watching and clicking.

The great distractor uses its power not merely to suck you in and keep you watching, but to shape new content, which in turn shapes us

YouTube is Edward Palmer's distractor at an almost imagineable scale. It knows what you watch, how long you watch it for, what you watched before and what you watched after and – given the pervasiveness of tracking tools across the wider internet – where you arrived from, who you are, and where you live. It knows this for billions of people, allowing it to sift out the slightest differences in tone and subject, differences which lose our attention or keep us locked in. YouTube is the perfection of a system of techniques held in common by *Sesame Street*, *Teletubbies*, and captology for the acquisition and retention of attention, stripped of all desire to *do* anything with that attention beyond profit from it. And this would be sad enough, were it not for the fact that such a system cannot be neutral. The great distractor uses its power not merely to suck you in and keep you watching, but to shape new content, which in turn shapes us.

One of the first signs that something was significantly amiss with this system came in 2016 when parents of very young children started complaining on internet forums and later in media outlets about the content their children were encountering on YouTube. Searching for popular kids television programmes such as *Peppa Pig*, *PAW Patrol*, and *Teletubbies*, children would land on – or more likely be directed to – knock-offs, parodies, and subversions of their favourites, often flavoured with violence, sex, and drug use.<sup>26</sup> Very young children – including toddlers and pre-schoolers – were being traumatised, their parents reported, and becoming frightened, inconsolable, and unable to sleep.<sup>27</sup> Google

spokespeople claimed each account was an isolated incident; they emphasised the amount of “good” content they hosted and the impossibility of full oversight; they argued, in effect, that there was nothing that they could do.

Digging into the vast quantities of children’s media on YouTube– as I began to do at the end of 2017 – revealed something even stranger. While it was indeed very easy to stumble upon or be led to disturbing and violent content aimed at kids – and the gravity of that should not be ignored – it was even easier to quickly lose oneself in a morass of content which seemed to have no meaning at all.<sup>28</sup>

Children’s television often has an element of oddness. I challenge anyone to watch James Earl Jones recite the alphabet in *Sesame Street*’s early episodes (available on YouTube) and not feel slightly threatened.<sup>29</sup> The *Teletubbies* episode “See-Saw”

was re-edited after parents reported that one sketch, featuring a cut-out lion chasing a cut-out bear under lowering skies was “unsettling”. Its focus on the needs of children often leads adults to remark on its incomprehensibility, as we saw with *Teletubbies*. But the algorithmically-engendered content that dominates and persists on YouTube is something else: uncanny, often creepy, and strangely, but unquestionably inhuman.




Fig. 6<sup>30</sup>

One aspect of this creepiness is the way in which human actors became meat for the algorithm. Take the example of *Bounce Patrol*, a popular channel with 16.4 million subscribers; a wholesome group of young childrens’ entertainers who post a couple of videos every week, and which rack up millions of views. The titles of these videos are revealing of the process at work: “Alphabet Christmas - ABC Christmas Song for Kids”, “Halloween Baby Shark | Kids Songs and Nursery Rhymes”, “Alphabet Halloween - ABC Halloween Song”, “Alphabet School - ABC School Song | Back to School”, “Alphabet Superheroes - ABC Superhero Song for Kids | Batman, Spiderman, PJ Masks, Incredibles, Hulk”, “Alphabet Transport - ABC Transportation Song for Kids | Vehicles, Phonics and Alphabet ABCs”, “Christmas Finger Family, Jingle Bells & more Christmas Songs for Kids! Popular Christmas Songs”, “12 Days of Christmas - Kids Christmas Songs |

Learn Counting for Kids | Popular Christmas Songs”, “Finger Family Song - Mega Collection part 3! Extended Family, Colors, Superheroes, Halloween & more” and so on and on and on. Infinite variations on the same themes, with titles designed to attract the attention of the algorithm, and containing songs and dances of the same remix and churn quality. We are a long way from *Sesame Street* – and this is the better end of the spectrum.

As certain videos become popular, more and more content creators try and fill the demand for the related search terms – a demand not from the viewers themselves (often, as we’ve noted, very young children), but from the recommendation algorithm. In response to the rising value of relevant search terms, videomakers rush to copy the content, resulting in an explosion of similar but subtly different versions of every single novelty hit, fad, and pre-school whim, from *Baby Shark* to the *Finger Family*. The result is the return of the Mechanical Turk: thousands of workers desperately impersonating machines, in the hope that they will do a good enough job to be chosen for promotion.

And that’s the good stuff. Most content is much, much worse. Instead of brightly-coloured children’s entertainers, the vast majority of the content created for children’s consumption on YouTube is churned out by anonymous animation sweatshops, or assembled automatically from cheap 3D assets and stock audio. Badly drawn versions of popular cartoon characters shudder to the sound of off-key nursery rhymes; counting games segue into toy commercials; the notions of “learning”, “education”, and “child development” are reduced to bare keywords for attracting advertising dollars.

At this level, the most disconcerting realisation is that it’s impossible to know exactly where this is coming from: who made it, where are they, and what was their intent? Are they individuals, pranksters, organised groups, corporate entities, or mere scripts, forgotten pieces of software still running out their instruction set and oblivious to the consequences? Here, at the bottom of the content barrel, we confront the most profound and urgent problem of these vast automated systems we find ourselves living within: how to attribute agency and responsibility for the forces which shape our lives?

To complete the circle – or perhaps the spiral – the failure to moderate this content in any meaningful way means that the streams collide. The automatically generated videos based on nursery rhymes meet the deliberate fakes made by parodists and trolls, resulting in a thick soup of disturbing and destabilising content

which continues to deepen: the algorithm turned in on itself, and the producers, moderators, viewers – all of us – abandoned to its increasingly odd demands.

And if only it was just the children. However, in recent years, it has become increasingly obvious that the factors at work here – the disavowal of responsibility for content, the pursuit of attention at all costs, the complexity of the systems of transmission and recommendation, and the obscurity of sources and origins – are having cataclysmic repercussions across society. It seems likely that rather than being a low point in the history of the internet, the now barely acknowledged scandal of children’s media that engulfed YouTube in 2018 will have been an early warning of far greater convulsions to come.

To take just one example, consider the rise of Flat Earth conspiracy theories. In the last decade, the incidence of belief in a flat earth has skyrocketed, and the culprit is overwhelmingly YouTube.

Among visitors at the first Flat Earth Conference held in Denver, Colorado in 2018, every single attendee questioned had come to the belief via watching videos not merely *shown* to them by YouTube, but actively *recommended* to them by the system.<sup>31</sup> YouTube had radicalised them, first by suggesting they watch a video about Flat Earth theories, and then

[T]o actually interfere with the money-making algorithm at the level required to actually protect children, democracy, and our sanity would mean reducing the scale of their business

continuing to suggest more and more similar videos, building up their belief that this was a well-supported and widely-held viewpoint. Moreover, as the YouTube-driven interest in Flat Earth conspiracies increased, the system encouraged more and more people – believers or not – to make Flat Earth videos, further building momentum.

The same process at work with the children’s videos recurs here. The algorithm can be compared to a traffic management system which is programmed to create optimal driving conditions. On seeing how people slow down when they pass a car crash, it gets the idea that people like car crashes, and so it sets out the conditions to create more of them. As amusing as this might be when it relates to a belief in a flat earth, it is less so when the worldview being disseminated is male supremacy, white nationalism, or religious fundamentalism, all of which can be found being pushed, in stronger or weaker forms, by YouTube’s automated systems.<sup>32</sup>

In fact, at the global scale, it doesn't entirely matter which particular horrific, hateful ideology or divisive culture war is being incited (which, again, is by no means intended to downplay their actual consequences). What each has in common is the distrust and division produced, whether it be in existing authorities, in common knowledge, in the institutions of science and democracy, in society at large or between one another. This is the process that is now being enacted by the systems we have chosen to live among. This is what underlies the confusing and conflicted world of the present.

When Google responded to reports of traumatised children on their platform by saying that there was nothing they could do about it, they were right, in a way. When they, and Facebook, and Twitter, and others, claim that the circumstances of electoral interference, the rise of extremisms and fundamentalisms, distrust in government and shared narratives, were beyond their control, they were right too, in a way.



Fig. 7<sup>33</sup>

There was nothing they could do; because to actually interfere with the money-making algorithm at the level required to actually protect children, democracy, and our sanity would mean reducing the scale of their business. For a company whose business *is* scale, which cannot operate *except* at scale, that presents an existential threat; treating the problem would mean they would cease to exist. And so there was nothing they could do. (Bowing to public pressure, Google has in the past couple of years taken some measures to address these issues, primarily in the form of limiting the revenue available for content aimed at children. Because this process is in itself automated – as it has to be, to operate at scale – it remains woefully insufficient to actually protect children from harm, but it has driven large numbers of legitimate producers off the platform.)

Ultimately, this is not an issue of technology, or of design. It is what happens to society when its most important functions – education, democratic debate, public information, mutual care and support – are wholly offshored to corporate

platforms dedicated to profit-seeking. There's no "healthy" way to educate children entirely within structures of capitalist exploitation. There's no way to entertain them safely either, or to perform any of the other tasks we undertake every day, like contributing to a shared public discourse, or reading the news, or trying to find reliable information, or meaningfully participating in democracy, when those tasks are entirely enfolded within those same structures. The story of the journey from *Sesame Street* to YouTube Kids is not one of advancing technology, but of receding care and diminishing responsibility. In this way we are all complicit, for continuing to support a system in which those who care the least are rewarded the most.

Edward Palmer and the Children's Television Workshop built the distractor in order to improve their real work. That work, as described by the *New Yorker's* film critic Renata Adler in 1972, was "to sell, by means of television, the rational, the humane, and the linear to little children."<sup>34</sup> Google and its ilk have instead concentrated on improving the distractor, and the result is the triumph of the irrational, the inhumane, and the non-linear. And it has made them very rich indeed.





## Endnotes

- 1 Teletubbies Dipsy And Lala. Pinclipart. <https://www.pinclipart.com/maxpin/iiRmooJ/>.
- 2 Swain, G. (1997). "TELETUBBIES: Are they harmless fun or bad for our children?; THE QUESTION ON EVERY PARENT'S LIPS." *The Mirror*, 23 May.
- 3 "Tele Tubby Mummy; Meet the woman every parent in Britain should be dreading... the kids' TV genius." (1997). *The Mirror*, 23 August.
- 4 Lepore, J. (2020). "How we got to Sesame Street." *New Yorker*, 11, May.
- 5 "Teletubbies" by Leo Reynolds is licensed under CC BY-NC-SA 2.0
- 6 Cooney, J. G., (1967). The Potential Uses of Television in Preschool Education. Children's Television Workshop, New York, NY.; Carnegie Corp. of New York.
- 7 Lesser, G. S. (1974). Children and Television: Lessons From Sesame Street. New York: Vintage Books.
- 8 Gladwell, M. (2000). The Tipping Point: How Little Things Can Make a Big Difference. New York: Little, Brown & Co.
- 9 "Vintage Ad #762: Can You Tell Me How to Play, How to Play With Sesame Street?" by jbcurio is licensed under CC BY 2.0.
- 10 Nelson, A. (2017). "I played Dispy in Teletubbies." *I News*, 31 March.
- 11 "Falwell Sees 'Gay' In a Teletubby." (1999). *The New York Times*, 11 February.
- 12 "How we made: Teletubbies." (2013). *The Guardian*, 3 June.
- 13 "Interview with a teletubby." (2010). *Network News*.
- 14 "Teletubbies" by christopheradams is licensed under CC BY 2.0
- 15 Fogg, B. J. (1998). "Persuasive computers: perspectives and research directions." Proceedings of the SIGCHI Conference on Human Factors in Computing Systems.
- 16 Fogg, B. J. (2003). Persuasive Technology: Using Computers to Change What We Think and Do. Morgan Kaufmann Publishers. ISBN 1-558-60643-2.
- 17 "Social Media Mix 3D Icons - Mix #2" by Visual Content is licensed under CC BY 2.0.
- 18 Leslie, I. (2016). "The scientists who make apps addictive." *The Economist*, 20 October.
- 19 "Stanford Class' Facebook Application Crosses 1 Million Installs." (2007). *TechCrunch*, 20 November.
- 20 Helft, M. (2011). "The Class That Built Apps, and Fortunes." *The New York Times*, 7 May.
- 21 Harris, T. (2013). "A Call to Minimize Distraction & Respect Users' Attention."
- 22 Stone, L. (2008). "Just Breathe: Building the case for Email Apnea." *Huffington Post*, 17 November.
- 23 Newton, C. (2018). "Google's new focus on well-being started five years ago with this presentation." *The Verge*, 10, May.
- 24 Clement, J. (2020). "Global logged-in YouTube viewers per month 2017-2019." *Statista*, 30, June.
- 25 Solsman, J. E. (2018). "YouTube's AI is the puppet master over most of what you watch." *CNET*, 10, January.
- 26 Palmer, A. and Griffiths, B. (2016).



“ASSAULT ‘N PEPPA Kids left traumatised after sick YouTube clips showing Peppa Pig characters with knives and guns appear on app for children.”  
*The Sun*, 10, July.

27 Maheshwari, S. (2017). “On YouTube Kids, Startling Videos Slip Past Filters.” *The New York Times*, 4, November.

28 Bridle, J. (2017). “Something is wrong on the internet.” *Medium*, 6, November.

29 “Sesame Street: James Earl Jones: Alphabet.” YouTube, uploaded by Sesame Street, 22 January 2010.

30 “youtube- kids telecharger sur portable” by downloadsource.fr is licensed under CC BY 2.0.


31 Landrum, A. R., Olshansky, A., and Richards, O. (2019). Differential susceptibility to misleading flat earth arguments on youtube. *Media Psychology*, DOI: 10.1080/15213269.2019.1669461.

32 Ribeiro, M. H., Ottoni, R., West, R., Almedia, V. A. F., and Meira, W. (2019). Auditing Radicalization Pathways on YouTube. arXiv:1908.08313v3 [cs.CY]

33 “SESAME STREET\*Google” by COG LOG LAB. is licensed under CC BY-NC-SA 2.0.

34 Adler, R. (1972). “Cookie, Oscar, Grover, Herry, Ernie, and Company: The invention of “Sesame Street.” *The New Yorker*, 27 May.





Designed by Yasmeen Safaie

---