

Florian Idenburg and LeeAnn Suen

SCAN

To be wary, to beware, or to be aware: there are choices to be made by the architect in handling the algorithm. Semantic hairsplitting might be something more than tedious or clever here: the persistence of inherited words and ham-fisted skeuomorphism reveals the elastic nature of cultural memory. Files, folders, and desktops, after all, still clutter office vocabularies and productivity software icons, even if all three reside exclusively on computers.

The christening of a class of objects as the “Internet of Things” tells us about something more than the linguistic conundrum of naming new technology. Awkward, ad hoc, vague, at times a touch graceless: our transition to life amidst the objects that collect and transmit data can be described in much the same terms as the phrase itself. In 1999, a brand manager at Procter & Gamble first referred to the Internet of Things in a presentation about using radio-frequency identification to track the flow of goods through its supply chain. The ambition was to create a real-time double of the unwieldy physical world in a format that could be visualized and managed remotely. A host of wireless communication, battery, and data storage technologies have since enabled formerly isolated, pre-programmed pieces of equipment to become connected and adaptive. We could, in retrospect, say that words failed us in that crucial first moment: these connected objects were dubbed “smart,” and thus their unendowed counterparts relegated to “dumb.” May linguists of the distant future choose not to mine these terms too deeply.

The design of mid-century white-collar work environments depended on a framework of typological distinctions and operational relationships between employees, offices, and equipment. Desks, chairs, computers, phones, staplers, calendars, mailboxes, paper, pens, coffee makers, routers, paper shredders, file cabinets, clocks, calendars, air conditioners, and thermostats were firmly placed into the category of “things,” distinct from “persons” and “places” in form, matter, agency, and relationship to the bottom line. As a corollary, “things” demanded a certain extent of monitoring, maintenance, and manipulation required by employees; they defined certain accommodations necessary in an office to be fit and used. In this way, these things defined work and they defined space. But the advent of smart equipment scrambles these distinctions and relationships: employees are less often manipulators, more often the manipulated. Indeed, machine learning algorithms programmed into smart devices are designed to progressively minimize manual operations. Again, linguistic inertia and our limited imaginations sheepishly reveal themselves. New devices inhabit old shells and signifiers in name and form, imbued though they are with algorithmic authority and expanded capacities for communication. Things have taken on the performative capacities of people and places but hide behind their unrelated precedents. Smart “thermostats” actively manage worker comfort and air quality; smart “calendars” keep time and serve reminders; smart “phones” envelop

This building makes money all night long.



Advertisement for Touch-Plate International, Inc., *Architectural Record*, April 1985, 196.

employees in virtual offices, no matter where they are.

This semantic misalignment is not trivial. Employees and offices are also flattened in this operation: they can also be made “smart” (and alas, they may also remain “dumb”). Employees tracked by quantifiable performance metrics can be made subject to remedial management in arenas heretofore unmeasurable and thus unmanaged. But while the goals of office management and algorithmic streamlining work in tandem, the design of offices is

another question entirely. Algorithms, like management, are driven and refined by the production and analysis of outcomes against initial questions of concern. Algorithmic design, despite its name, is more estranged than it is related to the practice of design done by humans.

In the late summer of 2018, WeWork was in negotiations to lease as many as twelve floors of office space in 1 World Trade Center, the most indignantly and insistently symbolic structure of the Manhattan skyline. The deal



Our intimacy with the inanimate, at home and work, embodied in an anemic array of smart technologies. From left: the Amazon Dash button, Google Nest, a bluetooth beacon, and Lumo Lift.


would have brought WeWork's Manhattan locations to fifty, moved the company into one of the most recognizable addresses in the borough, and knocked JPMorgan Chase out of the top spot as the largest tenant of commercial office space in the densest business district in the nation.

Hype surrounding the deal subsided as the deal fell through, but WeWork was undeterred. The company diplomatically insisted that they would persevere in pursuit of a lease to join the illustrious companies at the World Trade Center complex. By fall, the company signed leases on two other properties in decidedly less emphatic buildings nearer to Midtown. One building to the west, one building to the east: both unassuming constructions, marked by the terraced humility required by New York City's zoning regulations. But the building alone was never really the point: one WeWork is merely a node in a network of amenities, members, and naturally, other WeWorks.

WeWork branded a shared office leasing concept with a swagger and saturation made potent by a relentless pairing of the physical with the psychic: co-working looks like a mustard-colored, microfiber sofa with mid-century modern lines; innovation travels best on sculptural spiral staircases; collaboration breeds in dining table-sized groups under the fuzzy flocking of intumescent paint; community tastes like fruited water and local beer in mason jars; fulfillment can be attained through a quantifiable number of square feet of glass conference rooms. The buzz of productivity in a common work area is electrified by an urban touch of anonymity and tempered with the suburban security of exclusivity: the person sitting next to you isn't a coworker you know, but you're both

members, and you're one shoulder tap away from adding someone new, and knowable, to your network. WeWork aggressively curated an experience of knowledge work that tied a postmodern idea of replicable place to the Silicon Valley model of replicable results. At its peak in 2019, WeWork had deployed such a model with rigor across 5.2 million square feet in Manhattan alone.

What started as a side hustle fixing up old buildings in Brooklyn turned into a much different creature with the help of the 300-year plan and more than \$15 billion from its biggest backer, Masayoshi Son. The Japanese magnate's ultimate vision was end-to-end real estate development driven by artificial intelligence, and WeWork was just a first step. The company's *raison d'être* and competitive advantage lay in the breadth and depth of data generated from its members and its portfolio of "physical products." WeWork's close tracking of the relationship between design and usage of its shared office spaces, individually rented offices, and hundreds of thousands of desks yields analysis that minimizes the risk that normally accompanies investment in a workspace. This risk scales: freelance workers, sole proprietors, small companies, and large corporations each face risks proportional to their revenue generation and overhead costs and carry different abilities to absorb it. But WeWork grew both its physical products and data assets so it could absorb that risk for even the largest of Fortune 500 companies. While its brand and culture of membership was based on co-working, over one-fifth of WeWork enterprise customers were companies that manage over 500 employees. Corporations like IBM and Facebook signed on to have segments of their workforces be "Powered by We" in one way or another.



Scanners and sensors tag team to generate and maintain optimal lighting levels, cache, and productivity in WeWork's globally replicated workspaces. WeWorks from Busan to Buenos Aires stretch into the lightwood horizon.

WeWork's incremental acquisition of other data-mining and information management platforms, ranging from Fieldlens, a construction management platform, to Teem, a meeting room management system, demonstrates the company's ambitions for distilling all segments of workspace delivery into manageable, optimizable, and operational data. Research and development within the company delivered white papers on algorithmic test fitting and desk layout. Humans in the company were charged with focusing on "more high-value and complex" tasks, presumably including the monitoring and management of further data harvesting and analysis. WeWork's assets lay in the datasets and algorithms that, in concert, conducted the work of location scouting, space planning, thermostat setting, and supply ordering. WeWork's ability to capture and present the narrative of these workplaces through data allows it to propose a reduction in risk and time so great and so certain it's nearly impossible for any business to turn down. For every pairing of a physical object to a psychic motivator of work, there is also a value added paired to a risk reduced. WeWork made clear that any space could be optimally managed for work by circumventing questions of design entirely.

The company advertises its use of neural network models. By feeding the workings and preferences of existing locations and members into a database, WeWork teaches


its models to predict outcomes for locations, clients, and buildings that have yet to exist. Testing the probable usage of a particular conference room type in a hypothetical floor of a future WeWork in a particular location can be accomplished by a neural net in significantly less time and with significantly more accuracy than a human analyst, let alone a designer. Office space can be fitted out with sensors and infrastructure to create a closed loop of maintenance, to be as productive as the workers inside it. In being occupied, each office generates value for all the speculative offices that will be leased by WeWork in the future. The harvest is guaranteed to be full.

WeWork attempted to become a publicly traded company in 2019. In spectacular fashion, a wider investor evaluation of the risk that WeWork had taken on in amassing billions of dollars of long-term leases meant to be paid by short-term tenants caused a violent, collective raising of eyebrows. The close scrutiny that followed revealed gross negligence and self-dealing by the CEO and founder, Adam Neumann. His self-styled guru persona had been built up from superhuman quantities of hubris, venture capital, and visionary language melded onto a six-foot-five frame, and his control of the company had been cemented by his majority shareholder status. The breathless, retrospective analyses that followed the company's takeover by Son and Softbank focused as much on

Neumann's dope smoking and incompetent management as on any evaluation of the company's businesses. The company was swiftly reorganized. Two WeWork executives were brought up to replace the ousted Neumann; their first order of business was to remove the sauna and cold plunge pool from Neumann's office, and to reconfigure it to contain conference rooms and desks for WeWork employees.

manufacturing, transforming humans into tunable machines, required a myopic and unwavering understanding of the goal of human work.

Taylor's vision for measuring efficiency based on seconds worked and products moved per worker has been all but achieved and widely applied in manufacturing. But in the economy that trades in knowledge and creativity, Taylor's distant progeny—researchers in organizational behavior



The automated production of a WeWork office layout, as documented in Carl Anderson et al., "Augmented space planning: Using procedural generation to automate desk layouts," *International Journal of Architectural Computing* 16, no. 2 (2018): 164–177.


WeWork took on the task of tracking and optimizing the most efficient ways to manage, use, and lease office space. Rent is one of the largest operating expenses for any company; using data to better administer office spaces themselves seemed an alluring proposition. But the frameworks for evaluating such efficiency are not limited to tools and spaces: they are also put to work in evaluating the workers themselves. In 1880, Frederick Taylor didn't have software, but he did have an obsession with an idea of efficiency that was so rare at the time that it constituted a sort of sublime ideal. Companies did not yet even know to dream of the results that Taylor could forecast, then deliver. Taylor's magic was in knowing how to craft the process in order to achieve pre-calculated optimization goals based on isolated empirical study. His perspicacity in reimagining the role of the laborer in

like computer scientists Ben Waber and Alex Pentland from MIT—are attempting to break down the knowledge worker into traits that would have seemed irrelevant to Taylor, if only because they were once immeasurable: neuroticism, extroversion, openness, agreeability, and conscientiousness.

The stopwatch for neuroticism is a white plastic box with filleted edges, worn around the neck. The box is described and marketed by Waber's company Humanyze as a "smart" ID badge. In addition to a standard RFID tag, these sociometric badges are embedded with accelerometers, infrared transceivers, microphones, and Bluetooth. Their similarity to the widely implemented form of corporate security, the ID badge, is touted as one of its formal strengths: non-invasive and not alarming.

Non-invasiveness is noted as an important strategy to reduce barriers to compliance. But key functional considerations have required significant departures from its formal inspiration. The size of its battery and internal hardware results in a device with the thickness and heft closer to a smartphone than an ID card. Collecting data via the accelerometer and microphones requires the badge to be worn centrally, flat on the chest, and with two fixed points of lanyard contact to minimize unintentional swaying that would make noise in the data. And finally, the badge requires some conscientiousness to ensure that employees charge the badge and upload their data.

but also abandoned searches. Like WeWork's spaces, workers are evaluated not only by the actions that contribute directly to revenue, but by the value that can be harvested from their interactions within a network of other workers, communications streams, and objects. During the study period, employees tracked by Humanyze help their companies monetize the previously non-monetizable by making it visible. But unlike data about area, temperature, and light levels, the data on workers is analyzed for the subjective qualities of dominance in a conversation or a lack of ambition. And unlike remodeling a conference room or reprogramming a thermostat, the power dynamics involved in remediating a person are not



The Humanyze Badge, handily worn around the neck and flush to the torso, to collect anonymized employee biometrics and activity metadata. At the very least, it's less invasive than a microchip.


The badges record speech patterns, movement intensity, and physical proximity between badged employees and tagged office locations. Algorithms translate raw and relative quantities of frequency and distance into metrics like "team integration" and "employee exploration." These observations are presented alongside information harvested from "data exhaust": not only the emails sent, but also the emails left unsent; not only the files accessed,

able to be flattened or ignored.

Waber, Pentland, and their teams of researchers are dedicated to the development of observation methods that can replace surveys, self-reporting, and ethnographic observation to learn about how people work. The goal is complete information, not select data, and the infinite scalability of the product is paramount. The less models

need to extrapolate and correct, the more certain and the more valuable Humanyze's services can be. Personal tracking wearables can record, analyze, visualize, and return data back to a wearer in buzzes and beeps to suggest behavioral changes. In response to the COVID-19 pandemic, the sensor-based analytics platform Estimote, which was primarily focused on Bluetooth networks for retail, quickly released a workplace wearable badge that would buzz when workers violated the recommended distancing limits and made their movements recordable for future contact tracing. Researchers recognize that the greatest block to success for corporations still running on human labor are humans and their unpredictable, regulation-breaking natures.

malfunction can be judged against or alongside productivity in real time. Savvy companies will soon be able to close the loop in turning data into algorithmically-determined optimization actions, depending on what has been learned as the best solution for a particular shortcoming. While "radical transparency" policies at companies like Netflix and Bridgewater Associates have instituted "feedback" as an integral cultural experience and resource management process, the practice is still maintained and implemented by humans, and the result is brutal. One can imagine how an inherently machine-controlled system might be attractive to future companies, if only to minimize the number of tears shed on company property while optimizing productivity.



Humans, and the tiniest motes in the air between them, reduced to a diagram of trackers. Image courtesy of Estimote, 2020.

Humanyze's badges are an example of a workplace innovation that would bring humans into the Internet of Things, a milieu of workspace data once reserved for building management systems: heating and cooling systems, water pressure, fire protection, gas metering, security, and doors. Overall activity, interactivity, and

The term "feedback" first came into usage in the early twentieth century with the discovery that the output of an electrical amplifier, when partially fed back into the input, would increase amplification, but also cause the amplifier to screech and howl as the signal compounded. Future

linguists might again forgive the slippage between this systematic electrical effect and the human social and political relationships described by the same term: a performance evaluation passed from supervisor to employee; an open letter on a firm's hiring practices; an anonymous review on a website; a whistleblower's report against a CEO.

The ultimate implementation of surveillance and diagnostics relies on the maintenance of a smooth, lossless feedback loop: the fidelity of the sensing, transmission, and analysis of information derived from workplace activity by the algorithm, and the fidelity with which workers' behavior can be remediated according to a prognosis, re-sensed, and re-analyzed so quickly as to feel instant. Feedback necessitates this closed system—the inputs beget the outputs, which become the inputs, and on and on. Management might tweak data collection or behavioral feedback tactics based on evolving company goals, but the ideal loop excludes humans from the procedure entirely, self-managing and self-sustaining toward profit maximization and the ideal allocation of human resources. If total fidelity cannot be entrusted to humans, by removing humans, the work of managing a workplace might finally achieve perfect objectivity and efficiency.

The introduction of design questions into office management systems necessarily disrupts these ideal loops. Human design re-introduces human risk. Companies strive to make scanning, tracking, and other forms of management technology invisible by hiding them under banal glassy boxes or else innocuous by costuming them in friendly silicone shapes. The strategy opts for stealthy infiltration over conscientious consent to encourage compliance. But algorithms have no allegiance to any system of management, no inherent hunger for profit, and no prejudice for human character. Just as algorithms have every capacity to promote productivity, so too do they have the capacity to sow chaos, stimulate delight, and disrupt and redirect the power of a worker. Literacy and fluency in the design of such technologies, dispersed among a diversity of interests and motivations, will begin to reflect other desires in the office, and in time the Internet of Things will instead be subsumed into the wider world of humans.

X

This essay is an excerpt from the forthcoming *Human(s) Work: The office of good intentions* by Iwan Baan, Florian Idenburg, and LeeAnn Suen (Taschen, 2021).

Are Friends Electric? is a collaboration between e-flux Architecture and Moderna Museet within the context of

its exhibition *Mud Muses: A Rant about Technology*.

Florian Idenburg is a co-founder SO-IL, an architectural practice in New York. He is co-author of the book *Human(s) Work* (Taschen, 2022).

LeeAnn Suen is an architect based in Boston, Massachusetts. She is co-author of the forthcoming *Human(s) Work* (Taschen, 2021).

Are Friends Electric? is a collaboration between e-flux Architecture and Moderna Museet within the context of its exhibition *Mud Muses: A Rant about Technology*.