

Business Storytelling in the AI Revolution

David Michael Boje

Working paper for the World Scientific Encyclopedia of Business Storytelling Series 3

Sept. 30, 2026 7AM

Abstract

This essay proposes six volumes of Business Storytelling in the AI Revolution as a collective inquiry into how artificial intelligence reorganizes narrative authority, work, ethics, strategy, and human answerability. Its starting point is material rather than merely conceptual: an AI chip, held in the palm of the hand, is an historical condensation of mined materials, silicon wafers, lithography, global supply chains, data-center infrastructures, trained models, human annotation, and organizational decisions. The essay joins antenarrative, True Storytelling, socio-economic consultation, and sociotechnical systems theory to resist cognitive surrender to automated summaries, rankings, and claims of inevitability. It develops the Ghost Vortex as a way to name the layered traces of labor, data, design, and exclusion that remain active inside apparently seamless AI systems. The six volumes invite scholars and practitioners to document these relations, identify hidden costs, and co-design AI arrangements that enlarge rather than diminish human dignity, plural voice, and responsible action.

Keywords: artificial intelligence; business storytelling; Ghost Vortex; sociotechnical systems; SEAM

Introduction

The purpose of these six volumes is to reflect on the AI revolution as it is becoming, rather than only as executives, vendors, and our own habits of mind say that it ought to be. Our words, concepts, and biases are not innocent. They can make a moving field rigid. They can turn a question into a slogan, an opportunity into a foregone conclusion, and a human difficulty into a dashboard metric. I want this project to keep the field open long enough for readers, authors, workers, communities, and organizations to see what is being installed around them and inside them.

We are already in the worldly life of AI. It is present when a streaming service suggests the next program, when a map directs a driver, when a retailer forecasts demand, when a hospital routes a patient, when an airline app offers a fare, or when a university system ranks an application. These are not separate technological episodes. They are interwoven storytelling processes that select, rank, remember, omit, and make some possibilities easier to see than

others. The question is not whether business storytelling will be affected. It already is. The question is whose stories, bodies, labor, evidence, and moral answerability will count in the new arrangements.

The six proposed volumes do not treat AI as a single actor or a finished object. They approach it as a contested field of becoming. AI may augment judgment, help people recognize patterns, and make difficult work more accessible. It may also conceal the material and human conditions of its operation behind a smooth answer, a prompt window, or an executive promise. This essay therefore begins with a deceptively simple invitation: hold an AI chip in the palm of your hand. Ask what it is, how it came to be, who made it possible, what it consumes, and whose stories disappear when we call it merely intelligent.

The Chip in the Palm of the Hand

For a reader who knows nothing about chips, start here: a chip is not a tiny brain that appears by magic. It is a small piece of specially prepared silicon that has been patterned into a vast number of microscopic electronic switches. Those switches are called transistors. They can be arranged to store, move, and transform electrical signals. When many chips, memory modules, network links, cooling systems, and software are organized together, they can perform the calculations used to train and run an AI model. The chip is not the whole AI system. It is one material concentration within a much larger system of machines, people, data, energy, and decisions.

A contemporary AI accelerator makes the scale visible. NVIDIA reports that a Blackwell architecture GPU contains 208 billion transistors, built as two dies joined through a high-speed chip-to-chip interconnect and manufactured on a custom TSMC 4NP process (NVIDIA, 2026a). This is not a claim that every AI application uses Blackwell. It is a concrete example of the scale and integration now being pursued for AI computation. A semiconductor “node” is not the literal size of the chip. It names a manufacturing generation and a complex set of design and fabrication choices. The chip you hold is therefore already a history of scientific conventions, corporate investment, patents, laboratories, and long supply chains.

Its materiality begins long before the clean room. Quartz and other silicon-bearing materials are refined into extremely pure polysilicon; that material is grown into a crystal, sliced into wafers, polished, cleaned, and repeatedly exposed, etched, implanted, deposited, measured, and tested. Copper and other conductive metals carry signals through interconnect layers. Insulators and dielectric films keep some pathways apart. Chemical gases, photoresists, ultrapure

water, filtration systems, packaging substrates, and high-bandwidth memory are all part of the story. A chip is a layered artifact, not a single substance.

Lithography is the work of writing patterns onto those layers. At the leading edge, extreme ultraviolet lithography uses 13.5-nanometer light to print the most intricate layers; other layers continue to use deep ultraviolet systems (ASML, 2026). The technical description is almost surreal: a carbon-dioxide laser strikes droplets of tin to generate EUV light; multilayer mirrors guide that light because ordinary lenses do not work in this range; and the process must operate in high vacuum because EUV is absorbed by air. The wafer stage is positioned with extraordinary precision. What looks, in the palm of a hand, like a black or green rectangle is a physically exacting choreography of light, vacuum, chemistry, measurement, and repair.

The chip then enters another set of layers: advanced packaging, memory, servers, racks, switches, fiber, substations, backup generation, cooling equipment, buildings, and data centers. It also enters software layers: operating systems, drivers, compilers, model architectures, training code, safety filters, prompts, retrieval systems, and enterprise permissions. The International Energy Agency (2026) identifies AI and data centers among the dynamic sources of rising electricity demand. This does not mean that every data center uses the same amount of power or water, nor that one chip has one fixed footprint. It means that a serious account of AI cannot stop at the screen. The material question is also an energy, water, land, labor, and grid question.

Finally, the chip is not useful for generative AI without human and organizational work. Data are selected, licensed, scraped, cleaned, labeled, filtered, evaluated, and re-evaluated. Engineers tune models and systems; moderators and domain experts test harmful outputs; workers repair the errors that automated systems create; users supply feedback and new data through ordinary use. I do not repeat inflated claims about a fixed number of workers or rounds of training, because such counts vary by model and are not reliably disclosed. The better point is more basic: model training and deployment depend on continuing human labor and institutional judgment. The “machine” is never only a machine.

The Ghost Vortex and the Consent Machine

I call this layered afterlife the Ghost Vortex. It is not a mystical claim that a chip is haunted in the usual sense. It is a way of naming the traces that remain active after a system presents itself as smooth, neutral, and self-contained. The Ghost Vortex includes the training data whose origin is obscured, the worker whose task becomes invisible, the community asked to bear a new data center, the annotation rule that decides what is safe, the benchmark that becomes a

proxy for judgment, and the executive narrative that calls all of this inevitable. Like a palimpsest, the newest inscription does not entirely erase the marks beneath it.

The related idea of a Consent Machine asks how consent is made, narrowed, assumed, or withdrawn in these layered systems. A user may click “agree,” but that small gesture does not settle the consent of workers whose performance is monitored, authors whose work becomes training material, communities whose infrastructure is changed, or customers who cannot see why a decision was made. Consent is not exhausted by a checkbox. It is a continuing organizational practice of disclosure, contestation, refusal, repair, and answerability.

This is where business storytelling matters. Storytelling can either hide the Ghost Vortex or disclose it. A vendor story may promise frictionless intelligence while making workers responsible for the friction. A corporate story may call a ranking objective while withholding the data, goals, and exclusions that made it. A community story may name water, land, noise, and livelihood that a technical narrative leaves out. Antenarrative work keeps fragments in play before a dominant story seals them into a single authorized account (Boje, 2001). True Storytelling asks whether the account honors living relations rather than merely winning an argument (Larsen et al., 2021).

Theory and Evidence for the Six Volumes

The six volumes are not six silos. They are six faces of one changing field. Their relationship can be stated as a movement from material infrastructure and generative possibility, through authority and inclusion, to co-creation, strategic resistance, and socio-economic redesign.

Volume	Central question	Contribution to the Series
1 Generative AI Quantum Storytelling and Organizational Futures	How do generative systems participate in organizational sensemaking?	Develops quantum antenarrative, living machine storying, and accountable innovation.
2 Narrative Authority Cognitive Access and Institutional Legitimacy	Who is recognized as a knower when algorithmic systems classify and advise?	Examines authorship, disability, neurodiversity, provenance, surveillance, and institutional legitimacy.

Volume	Central question	Contribution to the Series
3 True Storytelling DEI and the Age of AI	Whose experience is amplified, stereotyped, or excluded?	Applies True Storytelling to dignity, representation, belonging, and plural voice.
4 Human AI Co Creation in Entrepreneurship and Strategy	When does AI augment rather than displace judgment?	Uses sociotechnical systems thinking and human systems integration.
5 Strategy Storytelling Against Cognitive Surrender	How can organizations resist passive acceptance of automated conclusions?	Uses strategic renewal and multi voiced analysis to question rankings, summaries, and financial narratives.
6 SEAM Socio Economic Consultation in the Age of AI	What hidden costs are created and how can affected people co design change?	Uses SEAM and qualimetrics to connect dysfunction, time, work design, and responsible redeployment.

Volume 1 Generative AI Quantum Storytelling and Organizational Futures

The first volume holds open the question of what it means for a generative system to participate in organizational sensemaking. It does not grant the machine a human soul, nor does it reduce it to a passive tool. It asks how people, models, data, artifacts, and institutions co-produce possibilities. The Ghost Vortex belongs here as a method of following the layered residues that corporate stories normally flatten.

Volume 2 Narrative Authority Cognitive Access and Institutional Legitimacy

The second volume asks who gets to speak, author, know, decide, and be believed. AI may assist accessibility and cognition, but it can also install new filters of legitimacy. If a model summarizes a student, worker, patient, scholar, or citizen, who can question the summary? What information, provenance, and appeal are available? This volume places disability, neurodiversity, authorship, and surveillance at the center rather than at the margin.

Volume 3 True Storytelling DEI and the Age of AI

The third volume brings True Storytelling to diversity, equity, inclusion, belonging, and representation. Its concern is not a decorative diversity statement attached to a technical system. It is whether the people represented by, evaluated by, and affected by AI can contest its categories and tell stories not already anticipated by the model. This is ethical resistance in the practical sense: preserving community, depth, somatic relation, and dignity against reduction to efficiency scores.

Volume 4 Human AI Co Creation in Entrepreneurship and Strategy

The fourth volume asks for the conditions of augmentation rather than displacement. Autonomous agents are increasingly described as systems that plan and execute multi-step work using models, tools, memory, and permissions (NVIDIA, 2026b). That description is useful, but it is not a governance model. Human–AI co-creation requires role clarity, meaningful human review, accessible evidence, error recovery, and the authority to stop or redirect the system. A sociotechnical design does not begin with “Where can we remove people?” It begins with “What human capability and collective learning should this arrangement enlarge?”

Volume 5 Strategy Storytelling Against Cognitive Surrender

The fifth volume names cognitive surrender: the habit of accepting a generated answer, ranking, forecast, or executive dashboard because it arrives quickly and appears mathematically composed. The remedy is not to ban calculation. It is to restore strategic storytelling as a discipline of asking whose assumptions are embedded, what evidence is missing, which futures were made unthinkable, and who has the standing to disagree. The volume can connect corporate finance, community prosperity, and ethical renewal without assuming that a model’s fluency is equivalent to wisdom.

Volume 6 SEAM Socio Economic Consultation in the Age of AI

The sixth volume develops socio-economic consultation as a method for AI transformation. SEAM and qualimetrics make visible the time, coordination failures, competence gaps, stress, rework, and lost knowledge that conventional business cases may not count (Savall & Zardet, 2008). The purpose is not only to criticize. It is to co-design bridge teams, training, work rules, and redeployment pathways with those who must live with the system. A hidden cost is a story waiting to be heard and translated into organizational action.

Agentic Commerce as an Illustrative Test Case

The notes supplied for this essay include a detailed account of a September 2026 “Muse” episode. I could not verify that event through reliable primary documentation, so I do not present it as a reported fact. The issue it raises, however, is real and timely. Agentic systems are being developed to carry out multi-step tasks, including shopping, customer service, planning, and other economic activity. Current accounts describe retailers exploring LLM-powered shopping and service agents, while technical descriptions emphasize the need for permissions, secure runtimes, policy controls, and human review (MIT Sloan Management Review, 2026; NVIDIA, 2026b).

For this reason, the supplied agentic-commerce table should be retained as an illustrative SEAM scenario, not as evidence about one named product. It maps four recurring sites of hidden cost: discovery and recommendation; checkout and processing; failure and dispute; and remediation and repair. At each stage, a smooth consumer story can transfer unseen time and risk to sellers, platform teams, support workers, banks, and customers. The appropriate research question is not “Does the agent work?” but “Who carries the labor of its success, its uncertainty, and its repair?”

A responsible research program would document actual system boundaries, agent identity, authorization, payment and credential handling, audit trails, appeal routes, and the distribution of liability. It would measure communication, coordination, and competence dysfunctions before and after deployment. It would then convene the people whose work and lives have been reorganized to co-author a redesign. This is a strong site of conversation among all six volumes: generative futures, authority, dignity, co-creation, strategic resistance, and SEAM hidden-cost analysis.

Implications for Authors Editors and Scholar Practitioners

The practical implication for the Series 3 workshop is that each proposed chapter should identify its AI object, its affected stakeholders, its evidence base, and its account of answerability. Authors can ask: What is the material infrastructure? What data and human labor are assumed? What story is being authorized? What voices are not yet heard? What visible and hidden costs are being transferred? What co-designed alternative is possible?

This is not a demand that every author become a semiconductor engineer or a data-center economist. It is an invitation not to let technical complexity become a reason for moral passivity.

The “chip in the palm” passage is a pedagogical doorway. It reminds us that a short answer on a screen is connected to a long chain of material and social relations. Once that relation is visible, scholarship can move from abstract alarm or abstract celebration toward careful, situated, and answerable inquiry.

The six volumes can therefore become a meeting place for different storytelling mountains: antenarrative, True Storytelling, Indigenous and together-storying traditions, sociotechnical theory, rhetoric, axiology, critical management studies, and socio-economic consultation. No one mountain has to erase the others. The task is to make their differences productive and to keep open the possibility that business storytelling in the AI revolution can be more plural, more materially aware, and more humane than the stories now being sold to us.

Conclusion

The AI revolution is not only a change in software. It is a change in the stories through which organizations perceive, decide, account, employ, and justify. Holding the chip in the palm of the hand interrupts the easy abstraction. It asks us to see silicon, light, water, electricity, work, data, and struggle where a finished product story sees only intelligence. The Ghost Vortex names the traces that remain. The Consent Machine asks who was permitted to say yes, no, or not yet. Series 3 offers six volumes in which scholars and practitioners can learn to see these relations together, resist cognitive surrender, and co-design arrangements in which AI serves rather than diminishes human answerability.

References

- ASML. (2026). EUV lithography systems. <https://www.asml.com/en/products/euv-lithography-systems>
- Boje, D. M. (2001). Narrative methods for organizational and communication research. Sage.
- International Energy Agency. (2026). Electricity 2026. <https://www.iea.org/reports/electricity-2026>
- Larsen, J., Boje, D. M., & Bruun, L. (2021). True storytelling: Seven principles for an ethical and sustainable change management. Routledge.
- MIT Sloan Management Review. (2026, February 18). Agentic AI, explained. <https://mitsloan.mit.edu/ideas-made-to-matter/agentic-ai-explained>
- National Institute of Standards and Technology. (2023). Artificial intelligence risk management framework AI RMF 1.0. <https://doi.org/10.6028/NIST.AI.100-1>

- NVIDIA. (2026a). NVIDIA Blackwell architecture. <https://www.nvidia.com/en-us/data-center/technologies/blackwell-architecture/>
- NVIDIA. (2026b). What are autonomous AI agents? <https://www.nvidia.com/en-us/glossary/ai-agents/>
- Savall, H., & Zardet, V. (2008). Mastering hidden costs and socio-economic performance. Information Age Publishing.
- UNESCO. (2021). Recommendation on the ethics of artificial intelligence. <https://unesdoc.unesco.org/ark:/48223/pf0000381137>