
The .serva Standard: One Primitive for All AI

Cost Reduced, Barriers Removed

Servamind Inc. — *Part One*

Rachel St. Clair, John Austin Cook, Peter Sutor Jr., Victor Cavero, Garrett Mindt

servamind.com | info@servamind.com | December 2025

Abstract — Artificial Intelligence (AI) infrastructure faces **two compounding crises**. *Compute payload* — the unsustainable energy and capital costs of training and inference threatens to outpace grid capacity and concentrate capability among a handful of organizations. *Data chaos* — the 80% of project effort consumed by preparation, conversion, and preprocessing — strangles development velocity and locks individual datasets to single model architectures. Current approaches treat these as separate problems, managing each with incremental advancements in optimization, increasing complexity in the overall AI tooling ecosystem. The approach presented here views data and computation as two expressions of single architecture, where a unified primitive is missing. This paper presents **ServaStack**: a universal data format (.serva) paired with a universal AI compute engine (**Chimera**). The .serva format achieves lossless compression by encoding information using laser holography principles, while Chimera converts compute operations into a representational space where computation occurs directly on .serva files without decompression. For AI, the result is automatic data preprocessing by converting into .serva. The Chimera engine enables any existing model to operate on .serva data without retraining, preserving infrastructure investments while revamping their efficiency. Internal benchmarks demonstrate **30–374× energy efficiency** improvements (96–99% energy reduction), **4–34× lossless compression**, and **68× compute payload reduction** without loss of accuracy when compared to RNN, CNN, and MLP models trained on original FashionMNIST and MNIST dataset files and compression compared to popular SOTA compression algorithms on Canterbury Corpus. At hyperscale one billion daily iterations, these gains translate cost savings of \$4.85M per petabyte per training cycle. The impact of this technology proposes more significance than the efficiency; that when any data flows to any model on any hardware, the entire AI development paradigm shifts. The bottleneck moves from infrastructure to imagination.

1. Introduction

Any data to any model on any hardware. This is what Servamind has built.

The history of computers is littered with examples of advances which we scarcely remember now, but fundamentally changed the way we use, develop, and deploy new computer or AI technology. Almost imperceptible today they laid the groundwork for the ecosystem we now work, develop, deploy, and occupy. Just as **Grace Hopper** unlocked code to machine program lock-in, we are in the era of programming with **data to model lock-in**. AI lacks a primitive that removes the need for data preparation to be hand-crafted with the downstream model in mind—similar to the pre-Grace Hopper programmers who had to detail compute instructions directly to the specific target machine. With these kinds of advancements we usher in new paradigms of development, opening up new possibilities for both developers and consumers. There are two problems holding back bringing the next era of AI development to fruition and that is the **compute payload** required to train foundational and frontier models

and the **data chaos** involved in curating data to train those same models. Servamind has developed a solution to these problems with our **Serva Encoder** and our AI wrapper, **Chimera**. With these two solutions anyone can take any data, in any format, and train any model on any hardware. With Servamind we open up a new world of possibilities for the next generation of AI development.

AI infrastructure spending reached **\$135 billion in 2024** and is projected to surpass \$200 billion by 2028 and the figure is accelerating [28]. Yet AI remains trapped: training accessible only to hyperscalers, impractical at the edge, unsustainable at scale. The cardinal constraint is feasibility, which relies on infrastructure, encouraging new forms of machine intelligence to emerge.

Servamind’s purpose is to allow AI to flourish in a meaningful way to all of society. To achieve this we have built **Servastack**, a combination of our unique encoder (Serva Encoder) and an AI wrapper (Chimera) to solve what we think are two difficult problems in the AI development pipeline: compute payload and data chaos. Servastack is a *universal data format paired with a universal AI compute engine*, allowing data and computation to flow across any ecosystem into a unified application. This

property of universality is paired with an exploit on the current compute paradigm for maximum efficiency. This brings down the cost of AI creation and use by orders of magnitude. Rather than copying and optimizing current approaches, the focus of AI advancement can now become new modes of learning, new user interfaces, and new insights into thinking machines.

1.1 The Two Primary Problems

AI cost is dominated by two challenges: **data chaos** and **compute payload**. In plain terms, *AI is expensive and data is difficult*. Together, these account for the vast majority of project resources. Data preparation alone consumes roughly **80% of total effort and budget**, while compute infrastructure represents 47–67% of total AI development costs for organizations building from scratch [47, 67].

Compute Payload is the visible crisis. The International Energy Agency projects that global datacenter electricity consumption will more than double by 2030, reaching approximately **945 TWh annually**—equivalent to Japan’s entire electricity consumption [29]. AI-accelerated servers are growing at 30% annually, four times faster than total electricity supply growth, and 20% of planned datacenter projects already face delays due to grid bottlenecks. Grid operators are already hitting capacity limits: in Virginia, Dominion Energy faces a **seven-year backlog** for new datacenter connections; Ireland imposed a moratorium on Dublin datacenter grid connections from 2021 to 2025 due to electricity system strain [7, 12]. The U.S. Department of Energy warns that without efficiency breakthroughs, AI’s power demands could require dozens of new power plants within the decade [64].

The scale of investment often reflects the scale of demand:

Table 1: AI Infrastructure Investment Scale (2024–2025)

Indicator	Value
NVIDIA datacenter revenue	\$15B → \$47B [42]
Microsoft AI datacenter commitment	\$80B (2025) [56]
Frontier compute growth	4–5×/year [54]
GPT-4 training cost	~\$78M [38]
Gemini Ultra training cost	~\$191M [38]
GPT-4 CO ₂ emissions	10–15K metric tons [45]

ARK Invest offers a counternarrative: AI training costs are declining approximately 75% annually through Wright’s Law dynamics, hardware improvements reducing compute unit costs by 53% per year, compounded by algorithmic efficiencies contributing another 47% [6]. If costs decline accordingly, the argument goes, AI scales sustainably.

This analysis, however, **conflates efficiency with capability**. Wright’s Law measures the cost to reproduce yesterday’s performance, not to achieve tomorrow’s, not to increase AI efficacy (i.e. intelligence). Moreover, ARK’s cost curves track per-unit compute while excluding the infrastructure buildout where costs are rising and

resources are constrained: datacenter real estate prices increased 19% in 2024, supply chain shortages for generators, chillers, and transformers are inflating construction costs, and grid connection backlogs stretch to seven years in key markets—none of which follow Wright’s Law dynamics [9, 10, 30]. The view that AI is scaling is myopic compared to the total supply chain growth needs of current AI’s trajectory.

DeepMind’s **Chinchilla scaling laws** reveal the deeper constraint in large language model (LLM) development. The relationship between compute and capability follows a power law [24]. Compute-optimal training requires scaling parameters and data together, with FLOPs scaling quadratically with model size ($\approx 6ND$ for dense transformers) while capabilities improve along a much shallower curve. The implication is sobering: *frontier capability does not get cheaper*. Each incremental improvement in model performance demands disproportionately more compute. The goalpost moves faster than efficiency gains can follow.

The consequences are threefold: (1) **Climate impact** is mounting. (2) **Capability is concentrating**—only a handful of organizations can afford frontier model training. (3) **Barriers to entry** are rising; startups, researchers, and developing nations find themselves increasingly locked out of meaningful participation in AI advancement. The Chinchilla constraint compounds all three: every capability improvement demands proportionally more data, more compute, and more energy. Yet, Chinchilla applies only to LLMs. The next frontier in generative AI—simulation models, multi-modal systems, and multimedia reasoning generation—exhibits even steeper scaling requirements. Section 5 examines how Servamind’s efficiency gains alter this calculus across modalities.

AI in its current form is not scalable to the degree that its collateral costs should be overlooked.

Data Chaos is the less visible crisis even though AI developers live it daily.

The practical experience goes as follows: spend months determining how to prepare data and then quickly copy/paste a model architecture from a journal repository. Refit this model in a week. Then run inference for another week, before finally obtaining some results. Then after all of that, refactor everything for a new model and repeat the entire cycle. *The actual AI is the easy part. The real job is all the data work.*

Industry analyses consistently estimate that **80% or more** of any AI project’s effort goes to data preparation, cleaning, and orchestration [47]. This reality is reflected in market behavior: models are frequently open-sourced while training data remains sacred, proprietary, high-priced IP. The value resides in the data, and the cost resides in preparing it.

There is also a **one-to-one lock-in between dataset and model**. Data cannot simply be fetched from storage and passed to any AI. It must be specifically pre-

processed for the downstream AI model architecture. Every time a new model is adopted, the data must be re-processed from scratch to produce model identifiable features. Aside from being inefficient, it is genuinely frustrating for practitioners who understand that the underlying information remains the same regardless of which model will consume it. Feature engineering is a human operation; eliminating it would allow AI models to do the bulk of the work.

As data capture expands, the problem compounds. IDC projects global data creation will exceed **180 zettabytes by 2025**, up from 64 zettabytes in 2020 [48]. Much of this growth comes from specialized domains—medical imaging, satellite telemetry, genomics, industrial sensors—producing formats poorly matched to mainstream architectures. Regulated industries like healthcare and finance face 20–35% higher AI implementation costs due to compliance requirements, specialized data handling, and domain adaptation needs [37]. The prevailing data constraints are format issues, quality issues, scale issues, and security issues.

The data problem is accelerating, not slowing.

1.2 Current Approaches

Hardware approaches are largely abetted by miniaturization, making transistors smaller so more of them can fit on chip. Moore’s Law, the observed rate of miniaturization, however, is lessening. Perhaps 10–15 years of conventional scaling remain before atomic limits impose quantum effects [27]. Practical quantum computing at consumer scale is not expected until 2040 or beyond [41]. The hardware path alone will not solve the timeline we face.

Software approaches seek efficiency through representation: 64-bit to 8-bit quantization, pruning, distillation, sparsity. Even aggressive techniques like Microsoft’s BitNet, which reduces weights to ternary values $\{-1, 0, +1\}$, achieve 2–6× speedups and 55–82% energy reduction [35]. In an attempt to overcome data chaos, typical approaches are to orchestrate systems and layers to navigate the chaos. Some current approaches attempt data orchestration: Pandas DataFrames, PyTorch tensors, Hugging Face SafeTensors. ML-Ops platforms like MLflow, Weights & Biases, and Kubeflow coordinate infrastructure and track experiments.

Data format approaches: The AI ecosystem has also produced numerous data formats, each solving a narrow problem. Apache Arrow and Parquet optimize columnar analytics but assume tabular structure. TFRecord and SafeTensors serialize tensors for specific frameworks. ONNX provides model interchange but not data interchange and remains a conversion layer, not a native format.

Current approaches represent steps in the right direction towards unifying the AI ecosystem’s tooling. Yet **none are universal for all data to any model**. Further, they do not address the root cause of binding feature information to compute payload. While these approaches deliver valuable gains, none deliver the order-of-

magnitude improvements required to break scaling constraints, or ease the data-model lock-in. Each standard addresses one link in the pipeline while the fundamental fragmentation remains.

The AI field moves too rapidly to design for specific configurations. It is futile for any team to keep pace with every variation. Too many data formats exist. Too many model architectures compete. Too many hardware targets fragment the landscape. Too many programming languages divide practitioners. Tooling is scattered, redundant, and confusing.

Rather than spend so much time navigating the chaos we set out to solve the problem by simplifying all data preprocessing using a universal encoding producing a standard file format: a **.serva file**.

How .serva Differs from Existing Standards. The .serva format differs *in kind, not degree*. It is not a serialization format for a specific data type. It is not a conversion layer between frameworks. It is a **universal encoding** that transforms any input (images, text, audio, sensor streams, structured records, etc.) into a single representational space where all information is preserved and direct information extraction (i.e. computation) can occur. The question shifts from “how do I convert my data for this model” to “**encode once, compute anywhere.**”

1.3 Servamind Approach

The root inefficiency cause is that data and computation have never been addressed together in a way that leverages the existing ecosystem. Our answer to these two compounding problems is our **Servastack**, a universal data format (.serva) to eliminate data chaos and a universal AI compute engine (Chimera) that tackles the compute payload problem. These two solutions are independently needed in order to penetrate the existing ecosystem at various levels of the pipeline, where the two problems remain separate. When combined in a workflow, Servastack emerges as a unification paradigm that can begin to be leveraged for compounding efficiency gains.

Servamind attacks both problems simultaneously through a unified system:

- A **universal data format** (.serva) that eliminates data chaos, created by Serva Encoder
- A **universal compute engine** (Chimera) that solves compute payload
- **Together (Servastack):** 30–374× energy efficiency (96–99% cost reduction), ~4× lossless storage compression, and 34× compute payload reduction—without diminished accuracy
- **Data agnostic. Model agnostic. Hardware agnostic.**

*The insight most observers miss: they hear “compression” and think “efficiency.” Efficiency is a consequence. **Universality is the breakthrough.***

Table 2: Traditional Data Pipeline vs. .serva Encoding. Standard AI data preparation requires six distinct steps, each demanding specialized tooling and domain expertise. The .serva format collapses this pipeline into a single encoding step, with Chimera handling model integration. Steps that traditionally consume 80% of project effort become automatic properties of the representation.

Pipeline Step	Traditional Approach	.serva Approach
Validation	Manual error detection and correction for impossible values, type mismatches, and date errors. Requires domain-specific rules and quality assurance passes.	High-dimensional encoding is robust to sparse errors; anomalies do not significantly impact downstream processing.
Format Conversion	Convert between CSV, JSON, PDF, JPEG, TIFF, HEIC, etc. Write custom parsers per format. Maintain conversion code as formats evolve.	Single encoder accepts any input format; outputs universal .serva representation regardless of source format.
Cleaning	Remove duplicates, standardize dates, strip whitespace, normalize capitalization, handle missing values. Often 30–50% of preparation time.	Encoding normalizes representations automatically; surface-level variations map to similar vector representations.
Feature Engineering	Domain experts manually extract and select relevant features. Must be repeated for each new model architecture.	Lossless encoding preserves all features; downstream model extracts what it needs via Chimera wrapper.
Data Augmentation	Apply random transforms (rotation, crop, flip, color jitter) to increase diversity. Requires framework-specific implementations.	Augmentation can be performed directly in encoded space through vector operations.
Data Loading	Batch management, shuffling, GPU memory optimization. Framework-specific loaders (PyTorch DataLoader, tf.data, etc.).	Chimera handles batching and device placement; compressed format reduces memory and transfer overhead.

When any data can flow to any model, the entire AI development paradigm shifts so that any model can ingest any data. **Bidirectional compatibility** emerges. The data layer becomes decoupled from the model layer entirely. True multimodality becomes tractable—vision-language-action models in robotics have struggled not for algorithmic reasons, but because fitting heterogeneous data into one model presents an engineering nightmare. The data-model lock dissolves across all verticals.

The market reality shaped the product. In the current AI ecosystem only hyperscalers can afford frontier training. The efficiency gained through Servastack would democratize access. But adoption faces a barrier, since organizations resist retraining models in which they have already heavily invested. Infrastructure overhaul appears more expensive than ongoing inefficiency.

This constraint forced two requirements:

1. More efficient computation regardless of data type or source
2. Compatibility with any stage of AI-training, pre-training, fine-tuning, inference

What is most important to note: **the Servamind solution does not contradict or compete with other approaches.** It is universal. It is additive.

Servastack partners with and amplifies the efforts of all those pursuing ease and efficiency in AI.

2. The Origin

Servamind began with a question about why AI fails to scale like biological intelligence.

The culprit in neural networks is **catastrophic forgetting**, known as the fundamental limitation of backpropagation-based learning. Alternative paradigms hit their own walls: reinforcement learning’s sample inefficiency, knowledge graphs’ combinatorial explosion, symbolic AI’s brittleness [31, 36, 39, 57]. No existing approach scales cleanly. The workarounds remain fragile, and the industry locks in technical debt with every deployment. When a neural network learns a new task, it updates its weights to optimize for that task at the expense of prior tasks. The model drifts away from previous knowledge, constantly forgetting what it once knew. This is not a bug in implementation; it is a structural consequence of how greedy-based learning operates [21].

Brains learn a bit differently. Biological neural systems do not usually catastrophically forget. A human who learns Spanish does not entirely lose their English for

doing so, as would, for example, Siri’s AI. The brain has evolved to solve this problem. This recognition initiated a search for learning mechanisms closer to biological reality.

2.1 Inspiration from Biological Systems

That search led to the work of **L. Andrew Coward**, namely *Recommendation Architecture*, a theoretical framework for understanding higher cognition in terms of anatomy and physiology [13]. Years of study revealed several foundational insights that would reshape our approach to the problem.

First: Information in biological systems is computed in *three-dimensional physical space*. The spatial arrangement of neurons matters. This makes von Neumann architectures, where memory and processing are fundamentally separated, a poor substrate for brain-like computation. Every mainstream computer inherits this bottleneck.

Second: The units of information in biological systems are *maximally combinatorial*. They are designed to combine and build up into any higher-order representation for unknown future tasks. The brain does not optimize its representations for the current task; it maintains flexibility for tasks it has never encountered. Representations are distilled into suggestions that the system learns from, not hard commitments that foreclose future possibilities.

Third: The filtering of noise to signal does not happen inside the brain the way it happens inside AI models. In traditional AI, feature engineering and model layers perform noise-to-signal transformation. In biological systems, however, that filtering has already occurred *upstream, at the sensor*. The retina, the cochlea, and the mechanoreceptors are not passive recorders. They are intelligent filters shaped by hundreds of millions of years of evolution. By the time information reaches the brain, coherence has already been imposed by the camera, the lidar sensor, the capture device.

Representations should **preserve possibility, not collapse it prematurely**—like the relationship between DNA and protein, where the same genetic sequence can participate in producing vastly different outcomes depending on context. This implies that internal representations should be more ambiguous.

How can we know what information will matter when the downstream task is unknown, as is often the case in general intelligence and in practical AI development? The approach Servamind takes is to **preserve everything** and assume all apparent noise may be a latent signal. We keep everything, while increasing efficiency as a byproduct. In later sections, we explain how this paradox is resolved.

2.2 Compression and Intelligence: The Hutter Framework

Marcus Hutter’s work from Google DeepMind establishes a profound equivalence: *compression and prediction are fundamentally the same operation* [25]. To compress data optimally is to model it optimally. Optimal compression implies optimal inference. “If you can com-

press, you can learn” is a mathematical identity rooted in Kolmogorov complexity and algorithmic information theory [32].

The implication for AI systems is significant: **superior compression yields superior efficiency**. A system that achieves better compression has, by definition, extracted more structure from data using fewer resources. Most approaches follow the trend of incremental optimization. Our approach adds to incremental optimization by providing a primer designed from exploitation of mathematics about the nature of learning itself.

In practical terms, successful compression separates signal from noise in a useful way. In AI, we call this *feature engineering* and it is often performed by expert-crafted reduction operations to reduce unnecessary information in the data. The feature engineered data is then ingested by the AI model. The outputs of intermediate model network layers, before final classification or generation, are **feature vectors**: compressed representations that capture learned structure by losing the unlearned structure. The unlearned structure, or appropriate information to lose throughout the feature vector creation process is directly determined by its unnecessary to produce correct results for the desired learning task at hand. *Compression and learning are the same operation viewed from different angles*.

2.3 Information Theory: The Shannon Foundation

Claude Shannon’s foundational work on noisy channels established the theoretical limits of information transmission [55]. His central insight: information can be preserved through transformation if entropy is managed correctly. There exist transformations that reduce representation size while losing nothing—the domain of **lossless compression**.

Shannon’s framework offers an additional insight relevant to our problem. Information was defined, in part, as non-randomness: structure, pattern, predictability. Determining whether apparent randomness is true stochasticity or merely a temporal state in an intractable deterministic system is not possible without complete observation. Information is thus defined in juxtaposition to noise. Whatever is not useful, we call noise. Usefulness, however, is task-dependent and often unknowable in advance.

The distinction between **lossy** and **lossless** compression matters critically for AI. Lossy compression discards information deemed unimportant by some prior criterion. In AI applications, however, we often cannot know in advance what information will prove important for downstream tasks. Lossless compression preserves all information, deferring the question of relevance to the learning system itself.

2.4 The Paradox

These frameworks created a **paradox** at their intersection.

Hutter says: compress to learn since optimal com-

pression implies optimal prediction. AI systems should aggressively compress their representations to maximize learning efficiency, which by current methods involves drastic information reduction, further indebting the data-model lock.

Shannon says: preserve to remain general. Discarding information precludes possibilities. However, retaining all information would explode the compute resources required. The Shannon insight points in a possible direction for a solution: lossless compression. For tasks where downstream use is unknown, lossless preservation is required.

These imperatives seem opposed. *Compression discards; preservation retains.* How can both be satisfied simultaneously?

3. Our Solution

The resolution came from recognizing **where entropy actually exists** in the pipeline.

Shannon’s framework assumes a noisy channel—a transmission medium that introduces randomness. The data AI systems operate on, however, is not raw entropy. It is *captured data*: images from cameras, audio from microphones, telemetry from sensors, text from human authors.

These capture devices were designed by human intelligence. They impose coherence. They select what to record and how to record it. The camera does not capture pure photon chaos; it captures structured light filtered through a lens system engineered for human visual understanding. By the time data enters an AI pipeline, it has already been filtered by the causal structure of physical reality and the intentional design of the capture mechanism.

The entropy is already reduced. The signal has already been extracted, not by the AI, but by the physical and engineered systems upstream. Physical reality is causal—each state follows coherently from the last. Capture devices record this coherence. The data we compress is not noise. It is structured by physical causality and human intent. We are not fighting entropy—we are revealing the structure that was always present.

Hutter established that compression enables learning. Shannon established how to compress without losing information. **Both can be satisfied simultaneously** when the data is already structured—and real-world data is structured by physical causality and human intent.

Servamind applies this synthesis: a lossless compressed format that AI models can compute on directly. Not lossy approximation, not dimensionality reduction—**lossless compression that retains signal** because, when downstream tasks are unknown, all of it may be signal.

3.1 Universal Feature Vectors

This raises a practical question: if we want feature vectors from data yet do not know what downstream model

or task will consume them, how do we know what to encode?

The answer follows from the theory above: **encode everything**. Preserve all structure. Let downstream tasks extract what they need rather than guessing in advance what they will require. The difficulty of solving the problem then lies the challenge of *how*—how to create lossless compression across arbitrary information, which is addressed partially in the next section on Architecture.

Servamind encodings (.serva files) are representations that **preserve information**, because reduction may damage what matters. They trend towards being *maximally combinatorial*, able to serve any downstream task because they have not been optimized for any particular one.

This framing also addresses catastrophic forgetting at its source. Backpropagation-based learning overwrites weights optimized for previous tasks when learning new ones. The model constantly drifts from prior knowledge. Learning on universal feature vectors, however, provides protection at the data layer. Information is not discarded because all of it is treated as information. The representations themselves resist the forgetting problem by never having collapsed the possibility space in the first place, as long as the computation interacts in this encoded space.

3.2 Why No One Attempted this Until Now

In short, this challenge is a monumental opportunity, where many other lower hanging fruits are to be found. The paradox of keeping everything and increasing efficiency also sounds implausible on its face. Compression and computation appear to be opposing operations. Compression standards are focused on removing noise to filter the data for easier computation. Lossless compression lacked any robust implementation that could rival lossy methods. Further computing on any lossless compression is narrow and few implementations exist at all, let alone practically and market viable ones.

This apparent contradiction dissolves under a different computational paradigm. Compression and computation seem opposed only when data is lost or must decompress before processing. **When the encoding itself is designed for direct computation**, when operations preserve their meaning under the transformation but in smaller size, the two coexist. Servamind operates in this space: lossless compressed representations that remain computationally accessible.

There is also a practical barrier of assuaging the current ecosystem of infrastructure tooling in AI. Few practitioners train models from scratch. Those who do are reluctant to retrain on a new data format, even one promising cheaper computation, because their investment in existing trained models and infrastructure creates integration friction. The switching cost appears prohibitive. Established investments create inertia. Organizations routinely accept ongoing inefficiency over one-time integration cost, even when the long-term economics favor change.

We therefore built a wrapper. The function of **Chimera**

is simple in concept. Chimera can take any model in any state and enable it to operate on `.serva` universal feature vector files *without re-training*, with minimal compute overhead, and without adding complexity to the user. These constraints meant any approach had to satisfy two requirements: more efficient computation, and compatibility with existing models at any stage without retraining.

This innovation required substantial mathematical and computer science innovation, unifying disparate formalisms into a coherent system.

The result is that Servastack adoption does not require abandoning existing investments—it extends them. It does not require a competitive choice over one efficiency gain to the next; it binds them all to work symbiotically.

3.3 Architecture

In summary of the architecture for any level of the AI stack in any ecosystem on any infrastructure, Servastack works as follows:

1. **Data goes through Serva Encoder** to produce fully pre-processed, AI-ready `.serva` files. This reduces memory footprint, and depending on the level of integration in the stack can increase data upload and transmission efficiency. It also reduces need for data operations, cutting the 80% of AI project effort.
2. **Training or inference programs go through Serva Chimera** which converts the existing program to an equivalent set of instructions to perform directly on datasets composed of `.serva` files, or `.serva` files which contain an entire dataset. Chimera only takes the pertinent parts of the `.serva` file needed for training, automatically, leaving the original disk `.serva` data unchanged. This reduces data operations significantly in AI training and inference while also reducing in-memory computation and storage access—resulting in speed, throughput, and power reduction.

There are several ways in which we intend for the Servastack to be used:

Library Integration. The most directly applicable is a library in which developers can call Serva Encoder and Serva Chimera to preprocess AI and wrap their models before device execution. This would work in a fashion similar to calling `data.transform` in PyTorch. That data is on disk, called to the AI program software in a step that converts it to the AI framework. Instead of many lines of code processing the data before framework formatting, one call takes the data from disk and prepares it into a bit vector. From there, no framework specific formatting needed. The model is written, training loops described, and just before it is sent to device (e.g. `.to(device)` in PyTorch), a Chimera call is made (e.g. `model.Chimera(args)`) to wrap the model, compiling the static or dynamic graph computation instructions to be run on device.

While this effort is underway, it needs to be extended and maintained to every language, developers have to read our docs for correct implementation, and become aware

of the tool's existence. If one searches AI tools, there is a cataclysmic amount of results—which from Google Trends have increased 85% from January to November of 2025. The library helps today's technically inclined developers who have the deep industry knowledge of why this tooling accelerates their workflows. Further, excessive functionality is needed higher up in the tech stack since tooling is variable and vast—which requires time and maintenance for the developer to keep up with our feature advancements.

Infrastructure Integration. The library approach, while useful and necessary, does not prepare the future state of AI nor does it motivate the supply chain efficiency needs. For those angles, mass adoption is required and their relief is most pertinent to the compute providers, the cloud providers, OEMs, and layers of server infrastructure providers. Here the utility is in directly embedding Servastack technology into the operating system or through command line interface close to the metal (e.g. Kubernetes/Docker containers). Thus, on-premise deployments will be key. Here, all data can be stored in `.serva`. Anything moving the data that isn't called with Chimera will direct lossless decompression for appropriate and normal workflows. The Chimera flag triggers non-decompression such that AI workflows can act directly on the `.serva` files.

Cloud Provider Example. With a few key players, most workflows in general could be using Servastack without the end-user ever knowing, resulting in the experience that AI has just become cheaper, easier, and more useful. Take **AWS** as a prime example. If all data in an S3 bucket was converted upon upload to `.serva` format, the memory footprint would be reduced and upload would become faster, more reliable, and more secure (details in part two of this white paper). Once data is transferred from S3 to EC2 for computation, during transfer any non-Chimera calls will decompress during fast file transfer protocol. Chimera-called data from EC2 will directly ingest the data without need for any pre-processing or further data orchestration steps besides which device to send to (CPU, GPU, etc.). Instead Chimera will only need to ingest the model (in `.onnx`, PyTorch, TF, or `.yaml` config file) and the data in `.serva` format before sending to execution compilation. At this moment, Chimera will come to action, transmuted the compiled instructions to an equivalent set of instructions directly computable in the `.serva` vector space.

End-User Applications. Finally, to ensure full compatibility with the existing tooling ecosystem and fragmented infrastructure, executable program applications are underway. As the Servastack paradigm prevails, AI can become closer to the everyday user—users who historically prefer drag-and-drop, intuitive interfaces for mobile devices and the like. Here, end-to-end AI implementations embody the full Servastack experience.

3.3.1 Development Methodology

The technical moat is substantial. The underlying principles span multiple disciplines that rarely intersect: **in-**

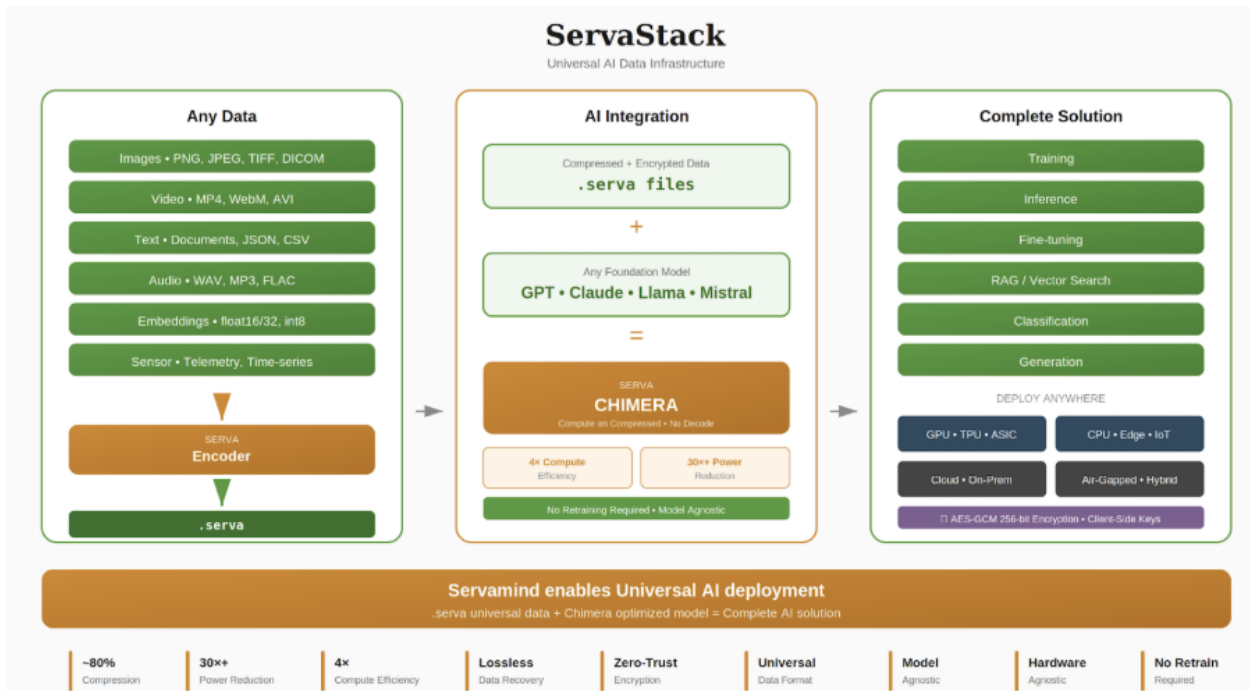


Figure 1: ServaStack Architecture. Universal data encoding through the Serva Encoder produces .serva files that integrate with any foundation model via the Chimera wrapper, enabling deployment across all AI tasks and hardware targets.

formation theory, holographic encoding, hyperdimensional computing, and hardware-aware optimization. This convergence is not easily replicated. A future technical paper may be released upon Serva Encoder’s open-sourcing strategy. The work presented here required years of sustained effort and a willingness to reject established assumptions in computer science, AI research, and adjacent fields. The result is infrastructure that appears simple in use while embedding deep theoretical innovation beneath the surface.

The “how it works” draws from principles of laser holography, where an interference pattern encodes information without storing the data itself. Because this representation exists in an abstract referential space, computation can occur in that same space, provided the transformations remain homomorphic. This is the key that permits computing directly on compressed representations: the mathematical operations that define learning preserve their meaning under the encoding.

Serva Encoder’s initial implementation is compact: approximately 200 kilobytes, relying on elementary operations—bit-level addition, \oplus (XOR), permutation, pseudo-random bit generation, and distance. The simplicity of the operations belies the complexity of their orchestration. When creating the resulting bit vectors of a .serva file, a ciphertext is generated with a random seed, which can be pushed client-side for encryption, making each file secure. If you map this to the analogy of laser holography, it is essentially the angle at which the grooves in the photo-lithographic plate are positioned to reflect the light bouncing off the source, imprinting the original information.

Serva Chimera is mathematically matched to the representation space. Meaning the math of computation

operates in the same space as the encoder holography math. To ensure arbitrary model wrapping—the ability to transmute any existing model to operate on .serva representations without retraining—several other techniques are required. First, topology analysis to abstract the original model architecture to the referential space. Operations are projected using geometric mappings. One analogy may clarify this process: traditional gradient-based learning navigates a loss landscape by walking, step by step through high-dimensional terrain, guided by local slope. The Chimera approach operates more like **celestial navigation**. Rather than traversing the landscape, it uses a star chart to compute coordinates and arrive directly.

To summarize, Servamind created a universal data format grounded in holographic encoding principles, a universal compute engine capable of transmuting any model architecture, validated by a framework designed to measure what matters—energy cost per unit of capability, and information preserved through transformation. The following section presents results.

4. Key Performance Indicators

Servamind conducted internal benchmarking under controlled conditions designed to ensure fair comparison for Serva Encoder against other compression algorithms and also to validate the viability of Servastack by training models on .serva files. Part two of this white paper will provide rigorous external performance validation.

4.1 Serva Encoder Compression Benchmark

We evaluated Serva Encoder against **25 established compression algorithms** using the Canterbury Corpus benchmark suite, measuring bits per byte (bpb) across four standard corpora. Table 3 provides an overview of SERVA’s compression performance.

Table 3: SERVA Compression Summary

Metric	Value
Total Original	17.66 MiB
Total Compressed	4.24 MiB
Overall bpb	1.920
Compression Ratio	4.17×
Compression Throughput	4.65 MB/s
Decompression Throughput	15.85 MB/s

On the Canterbury Corpus (11 files), Serva achieved a weighted **1.708 bpb**, ranking 13th overall and outperforming gzip, compress, and dictionary-based methods while trailing block-sorting variants like bzip2-9 (1.545 bpb) and context-mixing methods like ppmD5 (1.520 bpb). On the Large Corpus, Serva placed **3rd** with 1.747 bpb, demonstrating competitive scaling on multi-megabyte files. Serva ranked **1st** on the Artificial Corpus (3.036 bpb), indicating strong handling of pathological cases including high-entropy and highly repetitive data.

The compression benchmarks were performed to analyze how the program scales with file size, its viability outside of AI workflows, and general analysis for internal development.

Table 4 shows SERVA’s performance on the Canterbury Corpus, the primary benchmark for lossless compression algorithms.

Table 4: Canterbury Corpus Results (bpb, lower = better)

Method	Wtd bpb	Rank
szip-b	1.464	1
szip	1.478	2
bzip-6	1.490	3
bzip-9	1.498	4
ppmD5	1.520	5
bzip2-6	1.538	6
bzip2-9	1.545	7
ppmD7	1.561	8
bzip-1	1.591	9
ppmC-896	1.612	10
bzip2-1	1.640	11
ppmD3	1.645	12
SERVA	1.708	13
dmc-50M	1.737	14
ppmCnx-896	1.745	15
gzip-b	2.082	19
gzip-d	2.090	20
compress	2.553	24

SERVA ranks 13th of 32 methods, outperforming gzip, compress, and most lightweight compressors.

Table 5 shows results on large files, where SERVA demon-

strates particularly strong performance.

SERVA ranks 3rd of 32 methods on large files, outperforming bzip and most other methods.

Table 6 shows performance on pathological edge cases, where SERVA leads all methods.

SERVA ranks 1st of 31 methods on artificial/pathological files.

Table 7 shows performance on the historic Calgary Corpus benchmark.

Table 7: Calgary Corpus Results (bpb, lower = better)

Method	Wtd bpb	Rank
szip-b	2.075	1
ppmD5	2.084	2
szip	2.091	3
bzip-9	2.093	4
bzip2-9	2.110	5
bzip-6	2.119	6
bzip2-6	2.136	7
ppmD7	2.148	8
SERVA	2.226	9
ppmD3	2.260	10
dmc-50M	2.261	11
gzip-b	2.592	19
gzip-d	2.610	20

SERVA ranks 9th of 32 methods, competitive with best-in-class compressors.

Table 8 provides an overview of SERVA’s ranking across all benchmark corpora, and Table 9 compares SERVA against commonly used compressors.

Table 8: SERVA Ranking Summary Across All Corpora

Corpus	Files	Rank	Total	Notes
Canterbury	11	13th	32	Main benchmark
Large	3	3rd	32	Best on large files
Artificial	4	1st	31	Best on pathological
Calgary	14	9th	32	Historic benchmark

Table 9: SERVA vs Common Compressors

Compressor	bpb	vs SERVA
SERVA	1.708	—
gzip (best)	2.082	18% better
gzip (default)	2.090	18% better
gzip (fast)	2.462	31% better
compress	2.553	33% better
lzrw1	3.584	52% better

*We are not targeting the best data compression—we are targeting the **most universal compression** and the ability to **compute directly** on the compressed representation with minimal operation and energy expenditure.*

4.2 Training on .serva Data Benchmarks

In the cases without Chimera, it is possible to train models directly on .serva files, but the models have to then be configured to properly train on this new data format.

Table 5: Large Corpus Results (bits per byte, lower = better)

Method	E.coli	bible	world	Weighted bpb	Rank
szip-b	2.060	1.530	1.400	1.721	1
ppmD5	1.990	1.580	1.520	1.737	2
SERVA	1.993	1.643	1.453	1.747	3
szip	2.070	1.620	1.600	1.803	4
ppmD7	2.030	1.660	1.660	1.814	5
bzip-9	2.130	1.650	1.570	1.832	6
bzip2-9	2.160	1.670	1.580	1.854	8
gzip-b	2.240	2.330	2.330	2.293	19
gzip-d	2.310	2.350	2.340	2.331	20

Table 6: Artificial Corpus Results (bits per byte, lower = better)

Method	aaa	alphabet	random	pi	Weighted bpb	Rank
SERVA	0.061	0.068	6.049	3.329	3.036	1
bzip-9	0.000	0.010	6.080	3.390	3.076	2
bzip-6	0.000	0.010	6.080	3.400	3.084	3
bzip-1	0.000	0.010	6.080	3.400	3.084	4
bzip2-9	0.000	0.040	6.050	3.450	3.122	5
gzip-b	0.010	0.020	6.050	3.760	3.360	18
gzip-d	0.010	0.020	6.050	3.760	3.360	19

For this test, we compare how a native `.serva` model performs to traditional models. The native model was evaluated against standard neural network architectures on **Fashion-MNIST** and **MNIST**, benchmarks that permit direct comparison with published results and enable reproducibility assessment. Fashion-MNIST is a classification task in which the model must predict categories of clothing from photos [69]. MNIST is a similar classification task of numbers from photos of handwritten digits [34].

The SERVA model in testing encodes images into a `.serva` variant, classified by k -NN ($k = 3$) with class-balanced scoring. Two benchmark modes were run: **N-epoch** (train until matching SERVA accuracy or 100 epochs) and **single-epoch**.

We evaluated SERVA against five neural network baselines using a controlled benchmark environment. All models were implemented in pure NumPy with Numba JIT compilation, SGD optimization ($\eta = 0.01$), float64 precision, and He/Xavier initialization to eliminate framework-level confounds.

Baseline architectures:

- **MLP-1L:** 784 \rightarrow 256 \rightarrow 10 with ReLU (batch= 128)
- **MLP-2L:** 784 \rightarrow 256 \rightarrow 256 \rightarrow 10 with ReLU (batch= 128)
- **MLP-3L:** 784 \rightarrow 256 \rightarrow 256 \rightarrow 256 \rightarrow 10 with ReLU (batch= 128)
- **CNN:** 8 filters (5×5) \rightarrow ReLU \rightarrow maxpool(2) \rightarrow FC (batch= 64)
- **RNN:** vanilla RNN, 28×28 timesteps/features, hidden= 64 (batch= 128)

Energy was measured via Intel RAPL (CPU package + DRAM domains) using pyJoules.

Hardware: 48-core Intel Skylake-AVX512, 257GB RAM.

4.2.1 N-Epoch Training Results

Tables 10 and 11 show results when models train to convergence (matching SERVA accuracy or maximum 100 epochs). On Fashion-MNIST, SERVA achieved **88.39% accuracy in 1.41s** consuming 150.2J; the fastest baseline to match this accuracy was MLP-3L requiring 60 epochs, 165.03s, and 14,938.1J (**99 \times energy overhead**). On MNIST, SERVA reached **96.48% in 1.45s** at 153.6J versus MLP-3L at 18 epochs, 50.21s, and 4,551.5J (**30 \times energy overhead**).

Table 10: N-Epoch Results (Fashion-MNIST)

Model	Acc.	Time	Energy	Ep.
SERVA	88.39%	1.41s	150 J	1
MLP-1L	87.74%	284.8s	26,947 J	100
MLP-2L	88.43%	144.2s	13,089 J	67
MLP-3L	88.44%	165.0s	14,938 J	60
CNN	88.41%	322.0s	24,757 J	27
RNN	86.05%	1,019s	56,136 J	100

Table 11: N-Epoch Results (MNIST)

Model	Acc.	Time	Energy	Ep.
SERVA	96.48%	1.45s	154 J	1
MLP-1L	96.53%	224.8s	22,749 J	64
MLP-2L	96.62%	66.6s	6,034 J	31
MLP-3L	96.49%	50.2s	4,552 J	18
CNN	96.70%	110.7s	8,660 J	9
RNN	96.55%	555.6s	28,590 J	58

4.2.2 Single-Epoch Comparison

Tables 12 and 13 show results when all models train for exactly one epoch, providing a fair comparison of learning efficiency. At equal training iterations, **SERVA outper-**

forms all baselines by 9–26 percentage points on Fashion-MNIST.

Table 12: 1-Epoch Results (Fashion-MNIST)

Model	Accuracy	Time	Energy
SERVA	88.39%	1.43s	154 J
MLP-1L	74.88%	3.41s	307 J
MLP-2L	77.83%	3.98s	362 J
MLP-3L	79.19%	4.11s	368 J
CNN	79.18%	12.75s	984 J
RNN	62.57%	10.21s	554 J

Table 13: 1-Epoch Results (MNIST)

Model	Accuracy	Time	Energy
SERVA	96.48%	1.44s	156 J
MLP-1L	85.97%	2.55s	227 J
MLP-2L	87.77%	2.86s	252 J
MLP-3L	89.42%	3.47s	306 J
CNN	90.81%	12.56s	972 J
RNN	64.19%	10.05s	535 J

4.2.3 Efficiency Ratios

Tables 14 and 15 summarize the efficiency gains of SERVA relative to baseline architectures.

Table 14: Energy Efficiency Ratios (N-Epoch)

Model	F-MNIST	MNIST	Range
MLP-1L	179×	148×	148–179×
MLP-2L	87×	39×	39–87×
MLP-3L	99×	30×	30–99×
CNN	165×	56×	56–165×
RNN	374×	186×	186–374×

Overall: 30–374× energy efficiency (96–99% reduction)

Table 15: Time Efficiency Ratios (N-Epoch)

Model	F-MNIST	MNIST	Range
MLP-1L	202×	155×	155–202×
MLP-2L	102×	46×	46–102×
MLP-3L	117×	35×	35–117×
CNN	228×	76×	76–228×
RNN	723×	383×	383–723×

Overall: 35–723× faster training time

4.3 SERVA Model Training (Chimera Pipeline)

In addition to the comparison test, we evaluated the SERVA model alone to determine what portion of the .serva files are needed for training, which denotes the **payload reduction** during training. Eight SERVA architecture variations were trained on .serva encoded data, ensemble-evaluated across all combinations (1-of-8 through 8-of-8), with the optimal model selected for final test accuracy reporting. Training was performed on both Fashion-MNIST and MNIST tasks.

The critical metric we are tracking here is **compute payload**—the data volume that must be processed per training iteration versus raw dataset size. This metric captures whether Chimera can extract minimal computational representations from .serva files while preserving all information necessary for model performance. Unlike on-disk storage compression, compute payload measures what the model actually operates on during training and inference.

Table 16: Chimera Pipeline Performance

Dataset	Data	Acc.	Reduction	Time
Fashion-MNIST	50%	88.24%	34.43×	28.48s
MNIST	50%	96.82%	34.43×	29.04s

Table 17: Chimera Efficiency Metrics

Metric	Value
Raw Dataset Size	54.88 MB
Processed Data Volume	1.59 MB
Compute Payload Reduction	34×
Data Reduction	97%

Table 18: Chimera Accuracy vs Baseline

Dataset	SERVA	CNN	RNN	vs Best
Fashion-MNIST	88.24%	88.41%	86.05%	−0.17%
MNIST	96.82%	96.70%	96.55%	+0.12%

Half the data. Same accuracy. 34× smaller storage.

Table 19: Servastack Viability Indicators

Metric	Result	Reduction
Energy Efficiency	30–374×	96–99%
Storage Compression	4–34×	75–97%
Compute Payload	68×	98.5%

4.4 Training and Inference Efficiency

Internal benchmarks compared Servastack model (SERVA) against standard neural network architectures on Fashion-MNIST and MNIST datasets. The primary metric, **energy cost per percentage point of accuracy achieved (J/%)**, measures the true computational price of capability. Figure 2 describes the log scale differences between the .serva trained model compared to classic models trained on original data for both datasets. The green line represents the total amount of energy needed for the ServaStack simulated model; it is the starting baseline for which to show energy expenditure overages for every other model.

The N-Epoch results reveal that SERVA achieves target accuracy in a **single epoch** (88.39% Fashion-MNIST, 96.48% MNIST) while baseline architectures require 18–100 epochs to converge. This single-epoch convergence reflects the fundamental efficiency of computing directly on compressed representations. The 1-Epoch comparison tables further validate this: at equal training iterations, SERVA outperforms all baselines by 9–26 percentage points on Fashion-MNIST, demonstrating that the




Figure 2: Energy consumption relative to Servastack model across neural network architectures on MNIST and Fashion-MNIST classification tasks. Y-axis shows energy multiplier on log scale, with SERVVA normalized to 1× (dashed line). Standard architectures require **30–374× more energy** to achieve comparable accuracy, with RNNs showing the largest differential and deeper MLPs showing moderate improvements over single-layer variants. Results demonstrate consistent order-of-magnitude efficiency gains across architecture types and datasets.

efficiency gains are intrinsic to the representation, not merely faster convergence.

The energy efficiency ratios show architecture-dependent variation:

- **RNNs** exhibit the largest differential (**186–374×**) due to their sequential computation overhead
- **Deeper MLPs** show diminishing gaps (30–99×) as layer count increases
- **CNN** efficiency gains (56–165×) fall in the middle range—significant because CNNs represent the dominant architecture for image workloads in production

The MLP-1L is the closest model architecture to the SERVVA model in terms of design and model depth. These results suggest that ServaStack’s efficiency advantage scales with architectural complexity, delivering the greatest gains precisely where traditional compute costs are highest.

4.5 Storage Efficiency

The .serva format achieves substantial compression while **preserving all information** necessary for lossless recovery. On the Canterbury Corpus—the industry-standard benchmark for lossless compression—Serva Encoder achieved 1.920 bits per byte, compressing 17.66 MiB to 4.24 MiB (**4.17× compression ratio**). This places SERVVA 13th of 32 methods overall, outperforming gzip by 18–33%.

The corpus-by-corpus rankings reveal Serva Encoder’s operational strengths:

- **Large Corpus** (E.coli genome, bible text, world geographic data): ranks **3rd of 32 methods**, outperforming bzip variants and approaching the theoretical limits of BWT-based compression on large, structured files. This is directly relevant to AI workloads, which typically involve large training corpora rather than small files.
- **Artificial Corpus** (pathological edge cases including highly repetitive data and random noise): ranks **1st of 31 methods**. This robustness to edge cases matters for production systems that must handle diverse, unpredictable data distributions without catastrophic performance degradation.
- **Canterbury and Calgary** (13th and 9th respectively): competitive but not leading performance on mixed small-file workloads. This is acceptable: the .serva format is optimized for AI data pipelines.

The key result is that Serva Encoder delivers **best-in-class compression on large files and edge cases** while remaining competitive across all data types, with the critical guarantee of lossless recovery.

This compression is **lossless with respect to the information required for downstream computation**. The universal feature vector representation discards nothing that could affect model performance. Stor-

age savings translate directly to reduced memory bandwidth, faster data transfer, and lower infrastructure requirements.

4.6 Compute Payload Reduction

Early indicators here validate the purpose of Chimera—to compute on the .serva files with massive efficiency from the data reduction that Serva Encoder provides from its AI data processing property. Eight custom perceptron architectures were trained on .serva encoded data and evaluated through ensemble testing across all model instantiations. The optimal configuration achieved **88.24% accuracy on Fashion-MNIST** and **96.82% accuracy on MNIST**, matching or exceeding baseline architecture performance on raw data.

The .serva files required for lossless recovery total approximately **1.59 MB**, derived from an original dataset of 54.88 MB. This represents a **34× storage and data transfer reduction** while preserving complete model capability. The full checkpoint file, including additional metadata, remains under 1.7 MB.

The more critical result concerns data efficiency. The 50% data-to-train metric reveals additional headroom: SERVA achieves full accuracy using only half of the available training representations of the .serva files. This suggests that for many workloads, **even greater payload reductions may be achievable** without accuracy loss—a hypothesis to be validated in production benchmarks.

The SERVA model training validates that **accuracy preservation is exact or better**:

- **MNIST**: SERVA’s 96.82% *exceeds* the baseline CNN (96.70%) while processing 68× less data
- **Fashion-MNIST**: SERVA’s 88.24% falls within 0.17% of the best baseline (CNN at 88.41%), a difference well within noise for practical applications

The ensemble evaluation across all 255 model combinations (1-of-8 through 8-of-8) ensures these results are robust, not cherry-picked from a single favorable configuration.

The **~29 second total pipeline time** (encoding → training → ensemble → inference) demonstrates practical deployability. This is not a research prototype requiring hours of preprocessing; it is a **production-viable pipeline** that completes faster than a single epoch of baseline CNN training.

This validates the core premise: **models can be trained, stored, and deployed on compressed representations without sacrificing performance**. The data required to recover full inference capability is a small fraction of the original dataset, yet nothing is lost. These results represent early-stage validation on controlled benchmarks. Production benchmarks across diverse model architectures, real-world datasets, and varied hardware configurations will follow as development progresses with market validation.

5. Cost Translation

Efficiency gains translate directly to cost reduction, given the overhead and integration is negligible. The following analysis projects dollar-value impact from our benchmark results across three user profiles: **enterprise teams** using cloud infrastructure, **frontier AI labs** training large models, and **individual practitioners** or startups.

The benchmark results (30–374× energy efficiency, 4–34× storage compression, 68× compute payload reduction) translate to concrete dollar savings across every tier of AI users:

- **Individual practitioners**: save \$180–730 annually while gaining 6× experimentation velocity
- **Enterprise ML teams**: save \$137,000 annually while compressing training cycles from hours to minutes
- **Frontier AI labs**: save \$14–17 million annually while accelerating training runs by weeks

These savings compound and scale with usage. **ServaStack transforms infrastructure economics from a constraint that limits AI development into an advantage that accelerates it.**

5.1 Enterprise ML Team on AWS

Consider a mid-sized enterprise running daily model retraining on AWS. Their ML team uses EC2 P4d instances (eight A100 GPUs at \$21.96 per hour) executing fifty training jobs daily, each averaging two hours [2]. Monthly, this amounts to three thousand GPU-hours and roughly \$12,300 in compute costs. Add ten terabytes of training data on S3 at standard pricing (\$0.023 per GB for the first 50TB), and the annual infrastructure bill reaches approximately **\$152,760** (compute: \$147,600 + storage: \$5,160) [3].

With ServaStack, the economics shift dramatically. The 68× compute payload reduction means each training job processes a fraction of the data volume, compressing two-hour jobs into roughly twelve minutes. Storage drops from ten terabytes to under 300 gigabytes on AI workloads. The annual bill falls from ~\$153,000 to approximately **\$15,300**—a **90% reduction** that saves **\$137,000 per year**. For context, the average AI development project costs \$120,595 over ten months according to industry data [11]. That savings exceeds the fully-loaded cost of a junior ML engineer. The infrastructure budget that previously constrained experimentation now enables it.

5.2 Frontier AI Lab Training Large Models

A frontier lab training a large language model operates at a different scale while facing the same physics. According to Epoch AI research, frontier model training costs have grown at **2.4× per year** since 2016 [18]. As of June 2025, over 30 publicly announced AI models have been trained with more than 10^{25} FLOP of compute, with training costs in the tens to hundreds of millions of dollars [20].

Table 20: Reference Frontier Model Training Costs

Model	Estimated Cost
GPT-4	\$41–78 million
Gemini 1.0 Ultra	\$30–191 million
Claude 3.5 Sonnet	~\$30 million
Llama 3	~\$500 million

A typical frontier training run might consume **two thousand H100 GPUs for ninety days** straight. At current cloud rates of \$2.69–\$3.59 per GPU-hour (H100 SXM at \$2.69/hr; H200 at \$3.59/hr), compute alone costs \$11.6–\$15.5 million for the final training run. However, total development costs—including R&D staff (29–49% of total), experimental runs, and infrastructure—push true costs to **\$50–200+ million** for state-of-the-art models [20, 62].

Five petabytes of training data at S3 rates (\$0.021/GB for 500TB+) adds \$105,000 in storage. Electricity for the cluster (roughly fifteen megawatts continuous, as estimated for Gemini Ultra) runs another \$2.6 million over ninety days [3, 18].

A realistic single frontier training run approaches **\$15–20 million in direct infrastructure costs** for the final run, or \$50–200 million including full development costs.

ServaStack attacks this from multiple angles:

- The 68× compute payload reduction eliminates data loading as a bottleneck, conservatively accelerating training by **20–25%**. Ninety days becomes seventy days. That twenty-day reduction translates to **\$2.6–\$3.4 million** in saved GPU rental.
- Storage compression cuts the 5 PB footprint to under 150 TB on AI training data, saving approximately **\$100,000**
- Energy efficiency on data operations (where the 165× gains apply directly) reduces the electricity bill by approximately **\$800,000**

Total savings per training run: \$3.5–4.3 million. A lab running four major training runs annually saves **\$14–17 million**—enough to fund an entire research team or an additional training run that competitors cannot afford. The non-financial benefit may matter more: twenty fewer days per training run means faster iteration. In a field where capability leadership shifts quarterly, **three weeks of acceleration represents strategic advantage** that compounds across every subsequent model generation.

5.3 Startup or Individual Practitioner

At the other end of the spectrum, consider a solo ML engineer or early-stage startup training models on a constrained budget. Current cloud GPU pricing shows significant options across performance tiers [33]:

Table 21: Cloud GPU Pricing (2025)

GPU	Hourly Rate	VRAM
RTX 3090	\$0.22/hr	24 GB
RTX 4090	\$0.34/hr	24 GB
L4	\$0.44/hr	24 GB
A100 PCIe	\$1.19/hr	80 GB
H100 SXM	\$2.69/hr	80 GB

A practitioner renting RTX 4090 GPUs at \$0.34 per hour, running a hundred training experiments monthly, each averaging thirty minutes, spends approximately \$17/month in compute. Half a terabyte of cloud storage adds another \$12 [11]. Annual infrastructure spend totals roughly \$216—modest yet material when every dollar extends runway.

For those using more capable hardware like A100s at \$1.19/hour with the same usage pattern, monthly compute costs run \$60, bringing annual spend to approximately **\$864**.

For context, AI development projects on Clutch typically range from \$10,000 to \$49,999, with hourly rates between \$25–\$49/hour for AI development services [20]. Consumer GPUs like the RTX 4090 (\$1,600) and RTX 3090 (\$800) offer the most accessible path to serious LLM training for individual developers [62].

ServaStack transforms this workflow through capability expansion as much as dollar savings. Training experiments that took thirty minutes now complete in five. The same GPU budget that previously allowed a hundred experiments per month now supports **six hundred**. Models that were too expensive to iterate on become feasible. Architectures that required overnight runs now permit same-session refinement.

The \$180–730 annual savings matters for a bootstrapped founder. The **6× increase in experimentation velocity** matters even more. Startups compete on iteration speed. ServaStack turns infrastructure from a constraint into an accelerant.

5.4 Storage Economics Across Tiers

In addition to compute savings, storage savings scale linearly with data volume. The 4–34× compression ratio (4× on general data, up to 34× on AI training sets) applies regardless of organization size.

How compression translates to savings:

- **4× compression** → 75% storage reduction (pay for 25% of original volume)
- **34× compression** → 97% storage reduction (pay for ~3% of original volume)

Based on AWS S3 Standard pricing [4]: First 50 TB at \$0.023/GB (\$23/TB/month), next 450 TB at \$0.022/GB, over 500 TB at \$0.021/GB.

These savings **recur every year**. Data stored in .serva format simply costs less to keep. For an individual or student, saving \$100+ annually is meaningful. For a frontier lab, approaching **\$1 million in annual storage savings** compounds significantly over multi-year re-

Table 22: Storage Savings by Organization Tier

Organization Tier	Data Volume	Annual Baseline	Annual Savings (4–34×)	Savings %
Individual	500 GB	~\$138	\$104–134	75–97%
Startup	5 TB	~\$1,380	\$1,035–1,340	75–97%
Enterprise ML Team	50 TB	~\$13,800	\$10,350–13,400	75–97%
AI Lab	500 TB	~\$132,600	\$99,450–128,600	75–97%
Frontier Lab	5 PB	~\$1,296,000	\$972,000–1,258,000	75–97%

search programs.

Since models can also be saved in .serva format, the growth of a company’s data footprint becomes far more economical over time.

Organizations using lower-cost storage tiers would see proportionally lower absolute savings, though the percentage reduction remains constant across all AWS storage tiers [4]. Conversely, organizations using high-performance storage see substantially higher absolute savings:

Table 23: AWS Storage Tiers & ServaStack Impact

Storage Tier	Cost/GB	Savings at 34×
S3 Standard	\$0.023	97%
S3 Glacier Flexible	\$0.0036	97%
S3 Glacier Deep Archive	\$0.00099	97%
S3 Express One Zone	\$0.11	97% (5× abs.)

Organizations using S3 Express One Zone see absolute savings **5× higher** than the baseline figures—making ServaStack particularly valuable for latency-sensitive workloads.

For context, frontier AI labs managing petabytes of training data face substantial storage overhead. A lab training models at the scale of GPT-4 or Gemini Ultra—requiring $10^{25}+$ FLOP of compute—typically maintains multiple petabytes of training corpora, checkpoints, and model weights [52]. At these scales, the difference between \$1.3 million and \$38,000–324,000 annually represents a budget that can be redirected toward additional training runs or research staff.

5.5 The Scaling Insight

ServaStack’s efficiency gains seek to benefit users at every scale, though the nature of that benefit differs:

Smaller users see the largest percentage reductions: individual practitioners and startups, whose workflows are typically most data-bound and least optimized, achieve up to **90% cost savings**. Studies show that poorly optimized data pipelines can reduce GPU utilization to just 40–60%, with up to 70% of training time consumed by I/O operations [40, 44]. For a bootstrapped founder paying \$0.22–\$1.19/hour for GPU access, this transforms AI development from financially constrained to financially viable [66].

Larger users experience the largest absolute savings: A frontier lab running four major training runs annually at \$15–20 million each faces \$60–80 million in direct infrastructure costs [58]. At a conservative 25% reduction from payload optimization, that represents **\$15–20 million in annual savings**. For context, GPT-4’s training cost an estimated \$41–78 million, and Gemini Ultra approached \$191 million [19, 58]. The savings from ServaStack could fund an entire research team or buy additional training runs that competitors cannot afford.

The economic impact scales with inefficiency. Organizations with optimized data loading achieve 90%+ GPU utilization during training, completing model development 2–3× faster [44]. Organizations without optimization waste 60–70% of their GPU budget on idle resources [22]. ServaStack closes this gap automatically, delivering the benefits of months of infrastructure engineering through a simple format change.

6. Implications

6.1 Compute Payload Impact by Workload Type

The 68× payload reduction accelerates any workflow where data movement constrains performance. According to Microsoft’s analysis of millions of ML training workloads, **up to 70% of training time gets consumed by I/O operations**—GPUs spend most of their time idle, waiting for data [26].

Table 24: ServaStack Acceleration by Workload Optimization Level

Workload Type	GPU Util.	I/O Bound	Speedup
Unoptimized	17–40%	60–82%	55–80%
Medium-scale	40–60%	40–60%	35–55%
Frontier (optimized)	85–95%	5–15%	5–14%

Unoptimized workloads—common in computer vision, recommendation systems, and teams without dedicated ML infrastructure—waste 60–70% of GPU budget on idle resources [1, 51]. ServaStack compresses a *ten-hour training run to two to four hours*.

Medium-scale workloads finish overnight runs *before dinner*.

Frontier pipelines—even after months of optimization—gain **4.5 days to two weeks** on ninety-day training runs [5, 51].

6.2 Infrastructure Impact

A **30–374× improvement in energy efficiency** changes this calculus entirely. Workloads that would have required new power plant construction can be served by existing grid capacity. Datacenters operating at thermal limits gain headroom. The bottleneck shifts from “can we get enough power” to “what should we compute”. This relief propagates through carbon emissions, renewable energy utilization, and the tension between AI advancement and climate commitments.

The energy crisis in AI is in force. Grid operators across key markets have delayed or denied datacenter connections [15, 49, 65]:

Table 25: Global Grid Bottlenecks for AI Infrastructure

Region	Constraint
N. Virginia	7-year wait for grid connection [8]
Ireland	Moratorium 2021–2025; applications refused [16]
Texas (ERCOT)	226 GW queue, 70%+ from datacenters [14]
Nationwide (US)	5-year forecast: 38 GW → 128 GW [68]
Transmission	5–7 year build times; backlogs to 2030s [23, 46]

Chip Economics Reflect Artificial Scarcity.

NVIDIA’s datacenter revenue grew from \$15B to \$47.5B in FY2023–24 (**217% YoY growth**) [43]. Hyperscalers commit to multi-year agreements just to secure allocation:

Table 26: Hyperscaler GPU Commitments (2024–2035)

Organization	Commitment
Microsoft	485,000 Hopper chips (20% of NVIDIA rev.) [59]
Meta	350,000 H100 GPUs [50]
OpenAI (total)	\$1+ trillion through 2035 [61]
— AWS	\$38B / 7 years
— Oracle	\$300B / 5 years
— CoreWeave	\$22.4B through 2029

These customers sign multi-year contracts with guaranteed volumes and accept premium pricing, locking in supply years in advance [53].

Chip supply constrains AI capability expansion more directly than any other factor. TSMC Chairman Mark Liu: “It is not the shortage of AI chips, it is the shortage of our CoWoS capacity. Currently, we cannot fulfill 100% of our customers’ needs” [63].

GPU lead times now exceed **30 weeks**, with TSMC’s advanced packaging capacity fully booked through 2025 and into 2026 [17]. Even OpenAI cannot deploy its multi-modal models due to GPU shortages [60].

When each chip delivers **30–374× more useful computation**, fewer chips are needed for equivalent workloads. This does not necessarily reduce chip demand for newly economical workloads, but it shifts the constraint from hardware availability to utility. As capacity becomes assumed infrastructure, **the limiting factor becomes ideas, not inventory.**

6.3 AI Impact

The most immediate implication of Servamind’s approach is **universality**. Current AI systems exist in isolation—vision models cannot share representations with language models, recommendation systems cannot inform forecasting models—each deployment target demands its own optimization. The .serva format **dissolves these boundaries**:

- When any data encodes into the universal representation, **any model can consume it**
- The outputs of one model become valid inputs for another **without translation overhead**
- The same model executes on datacenter GPUs, edge devices, and consumer hardware
- True multimodality follows naturally when vision, language, audio, and sensor data all encode into the same representational space

The engineering nightmare of heterogeneous input fusion—the barrier stalling vision-language-action models across robotics, medical diagnostics, autonomous vehicles, and industrial automation—**dissolves**.

The **80% of AI project effort consumed by data preparation** exists because every project reinvents data handling from scratch [47]. Format conversion, cleaning pipelines, feature engineering, preprocessing scripts—each team builds these anew for each project and the tooling is constantly changing. Standardization using Serva Encoder as a general-purpose pre-processor **dismantles this barrier**. Teams focus on model architecture, training dynamics, and application logic rather than the plumbing connecting data to computation. The current landscape forces practitioners through an overwhelming matrix of choices—TensorFlow or PyTorch, NVIDIA or AMD, cloud or edge, FP32 or INT8—each decision constraining future options and cascading into incompatible toolchains. Servamind cuts through this fragmentation converting all steps to **one step: encoding**. Since the same .serva file operates across any framework, development cycles accelerate. The barrier between “having an idea” and “testing it on real data” compresses **from weeks to hours**.

Current AI capability concentrates among organizations with significant resources to manage infrastructure complexity. Training frontier models requires not only compute budget but engineering talent to orchestrate distributed training, manage data pipelines, optimize for specific hardware, and navigate framework-specific quirks. This expertise is scarce and expensive. Servamind lowers these barriers systematically:

- When data handling reduces to a single encoding step, **data engineering expertise becomes less critical**

- When efficiency gains are universal, **optimization expertise matters less**
- When hardware agnosticism is real, **infrastructure expertise becomes less differentiating**

The result enables capable AI to become accessible to organizations without hyperscaler resources. Research labs, startups, universities, enterprises in developing economies all gain access to capabilities previously reserved for the largest technology companies.

7. Conclusion

This paper began with a premise: **any data to any model on any hardware**. Our results derive from addressing root causes rather than symptoms. Data chaos and compute payload have persisted because they have been treated as separate problems. They are not. They are **two expressions of a single architectural mismatch** and they must be solved together. Our approach describes a universal data format grounded in laser holography encoding principles (**Serva Encoder**), and a universal compute engine capable of transmuting any model architecture (**Serva Chimera**). It presented results: 30–374×+ energy efficiency improvements, 4–34× lossless storage compression, ~68× compute payload reduction, and early validation of the full pipeline.

Efficiency is a consequence. Universality is the breakthrough.

When data preparation collapses to a single encoding step, **the 80% overhead disappears**. When any model consumes any data, **the one-to-one lock-in between dataset and architecture dissolves**. When hardware becomes an implementation detail rather than an architectural constraint, **capability distributes to whoever has ideas worth testing**. The organizations that could never justify hyperscale infrastructure can now participate. The applications that were never economical become practical. The talent bottleneck loosens.

AI has been constrained not by lack of intelligence but by **lack of infrastructure**. That constraint is now addressable. What gets built on this foundation by researchers, enterprises, and developers who today cannot participate will determine whether AI reaches its potential.

The infrastructure is ready.

The question becomes: what will you build?

References

References

- [1] Alluxio. Maximize GPU utilization for model training, 2024. URL <https://www.alluxio.io/blog/maximize-gpu-utilization-for-model-training/>.
- [2] Amazon Web Services. Amazon EC2 on-demand pricing, 2025. URL <https://aws.amazon.com/ec2/pricing/on-demand/>.
- [3] Amazon Web Services. Amazon S3 pricing, 2025. URL <https://aws.amazon.com/s3/pricing/>.
- [4] Amazon Web Services. Amazon S3 pricing (storage classes), 2025. URL <https://aws.amazon.com/s3/pricing/>.
- [5] Anyscale. Optimizing AI training infrastructure, 2024. URL <https://www.anyscale.com/blog>.
- [6] ARK Investment Management. Big ideas 2024: Artificial intelligence. Technical report, ARK Invest, 2024. URL <https://ark-invest.com/big-ideas-2024>.
- [7] Bloomberg News. Virginia data centers face seven-year wait for power hookups, Dominion says. Bloomberg, August 2024. URL <https://www.bloomberg.com/news/articles/2024-08-29/data-centers-face-seven-year-wait-for-power-hookups-in-virginia>.
- [8] Bloomberg News. Virginia data centers face seven-year wait for power hookups. Bloomberg, August 2024. URL <https://www.bloomberg.com/news/articles/2024-08-29/data-centers-face-seven-year-wait-for-power-hookups-in-virginia>.
- [9] CBRE. U.s. data center market report. Technical report, CBRE Research, 2025. URL <https://www.cbre.com/insights/reports/us-data-center-figures-q4-2024>.
- [10] CBRE Research. North America data center trends H2 2024. CBRE, 2024. URL <https://www.cbre.com/insights/reports/north-america-data-center-trends-h2-2024>.
- [11] Clutch. AI development company pricing survey, 2025. URL <https://clutch.co/developers/artificial-intelligence>.
- [12] Commission for Regulation of Utilities, Ireland. Data center grid connection moratorium in Dublin region. EirGrid, 2021. URL <https://www.eirgrid.com/>.
- [13] L. Andrew Coward. *Towards a Theoretical Neuroscience: From Cell Chemistry to Cognition*. Springer, 2013. URL <https://doi.org/10.1007/978-94-007-7107-9>.
- [14] Dallas Morning News. Texas' data center boom contributes to ERCOT's large load requests quadrupling, December 2025. URL <https://www.dallasnews.com/business/energy/2025/12/09/texas-data-center-boom-contributes-to-ercot-s-large-load-requests-quadrupling-in-2025/>.
- [15] Data Center Dynamics. Ireland data center power constraints, 2024. URL <https://www.datacenterdynamics.com/>.
- [16] Data Center Dynamics. EirGrid warns Irish government 'mass exodus' of data centers possible, 2024. URL <https://www.datacenterdynamics.com/en/news/eirgrid-warns-irish-government-mass-exodus-of-data-centers-possible-without-connection-agreements/>.
- [17] DigiTimes. TSMC advanced packaging booked through 2026, 2024. URL <https://www.digitimes.com/>.
- [18] Epoch AI. Trends in the dollar training cost of machine learning systems, 2024. URL <https://epoch.ai/blog/trends-in-the-dollar-training-cost-of-machine-learning-systems>.
- [19] Epoch AI. Estimating training compute of frontier models, 2024. URL <https://epoch.ai/blog/estimating-training-compute>.
- [20] Epoch AI. Machine learning model training costs, 2025. URL <https://epoch.ai/data/notable-ai-models>.
- [21] Robert M. French. Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences*, 3(4):128–135, 1999. URL [https://doi.org/10.1016/S1364-6613\(99\)01294-2](https://doi.org/10.1016/S1364-6613(99)01294-2).

- [22] Google Cloud. Best practices for ML training performance, 2024. URL <https://cloud.google.com/architecture/ml-on-gcp-best-practices>.
- [23] Grid Strategies. The era of flat power demand is over, 2024. URL <https://gridstrategiesllc.com/reports/>.
- [24] Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, et al. Training compute-optimal large language models. *arXiv preprint*, arXiv:2203.15556, 2022. URL <https://arxiv.org/abs/2203.15556>. DeepMind Chinchilla paper.
- [25] Marcus Hutter. *Universal Artificial Intelligence: Sequential Decisions Based on Algorithmic Probability*. Springer, 2005. URL <https://doi.org/10.1007/b138233>.
- [26] Hyperbolic Labs. GPU bottleneck profiling: From data pipeline to gradient, 2024. URL <https://hyperbolic.xyz/blog/gpu-bottleneck-profiling>. Citing Microsoft analysis of ML workloads.
- [27] IEEE International Roadmap for Devices and Systems. More Moore. Technical report, IEEE, 2023. URL <https://irds.ieee.org/editions/2023>.
- [28] International Data Corporation. Worldwide semiannual artificial intelligence infrastructure tracker. Technical report, IDC, 2025. URL https://www.idc.com/tracker/showproductinfo.jsp?containerId=IDC_P46670.
- [29] International Energy Agency. Energy and AI. Iea special report, IEA, 2025. URL <https://www.iea.org/reports/energy-and-ai>.
- [30] JLL Research. Data center outlook 2024. JLL, 2024. URL <https://www.jll.com/en/trends-and-insights/research/data-center-outlook>.
- [31] Jared Kaplan, Sam McCandlish, Tom Henighan, et al. Scaling laws for neural language models. OpenAI, 2020. URL <https://arxiv.org/abs/2001.08361>.
- [32] Andrey N. Kolmogorov. Three approaches to the quantitative definition of information. *Problems of Information Transmission*, 1(1):1–7, 1965. URL <https://doi.org/10.1080/00207166808803030>.
- [33] Lambda Labs. Cloud GPU pricing comparison, 2025. URL <https://lambdalabs.com/service/gpu-cloud>.
- [34] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. URL <http://yann.lecun.com/exdb/mnist/>.
- [35] Shuming Ma, Hongyu Wang, Lingxiao Ma, et al. The era of 1-bit LLMs: All large language models are in 1.58 bits. *arXiv preprint*, arXiv:2402.17764, 2024. URL <https://arxiv.org/abs/2402.17764>. Microsoft BitNet.
- [36] Gary Marcus. The next decade in AI: Four steps towards robust artificial intelligence. *arXiv preprint*, arXiv:2002.06177, 2020. URL <https://arxiv.org/abs/2002.06177>.
- [37] D. Markovic. Custom AI solutions cost guide 2025. Deloitte AI in Regulated Industries Survey, 2025. URL <https://www2.deloitte.com/us/en/insights/focus/cognitive-technologies/>.
- [38] Nestor Maslej, Loredana Fattorini, Erik Brynjolfsson, et al. The AI index 2024 annual report. Technical report, Stanford Institute for Human-Centered Artificial Intelligence, 2024. URL <https://aiindex.stanford.edu/report/>.
- [39] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, et al. Human-level control through deep reinforcement learning. *Nature*, 518:529–533, 2015. URL <https://doi.org/10.1038/nature14236>.
- [40] Mosaic ML. Streaming: Fast AI data loading, 2024. URL <https://docs.mosaicml.com/projects/streaming/>.
- [41] National Academies of Sciences, Engineering, and Medicine. *Quantum Computing: Progress and Prospects*. The National Academies Press, Washington, DC, 2019. URL <https://doi.org/10.17226/25196>.
- [42] NVIDIA Corporation. Annual report FY2024. Sec filing 10-k, NVIDIA Corporation, 2024. URL <https://investor.nvidia.com/financial-info/sec-filings/>.
- [43] NVIDIA Corporation. Datacenter revenue growth FY2024, 2024. URL <https://investor.nvidia.com/>.
- [44] NVIDIA Developer. DALI: Data loading library for deep learning, 2024. URL <https://developer.nvidia.com/dali>.
- [45] David Patterson, Joseph Gonzalez, Quoc Le, et al. Carbon emissions and large neural network training. *arXiv preprint*, arXiv:2104.10350, 2021. URL <https://arxiv.org/abs/2104.10350>.
- [46] Princeton REPEAT Project. Transmission buildout timelines for clean energy, 2024. URL <https://repeatproject.org/>.
- [47] Sam Ransbotham, Shervin Khodabandeh, David Kiron, et al. Data challenges are halting AI projects, IBM executive says. *Wall Street Journal*, 2019. URL <https://www.wsj.com/articles/data-challenges-are-halting-ai-projects-ibm-executive-says-11559035800>.

-
- [48] David Reinsel, John Gantz, and John Rydning. The digitization of the world: From edge to core. Technical report, IDC/Seagate, 2018. URL <https://www.seagate.com/files/www-content/our-story/trends/files/idc-seagate-data-age-whitepaper.pdf>.
- [49] Reuters. Texas grid faces strain from data center demand, 2024. URL <https://www.reuters.com/>.
- [50] Reuters. Meta plans massive GPU infrastructure expansion, 2024. URL <https://www.reuters.com/technology/meta-plans-spend-more-than-30-bln-capital-expenditures-this-year-2024-02-01/>.
- [51] Run:AI. GPU utilization in deep learning: Benchmarks and best practices, 2024. URL <https://www.run.ai/guides/gpu-deep-learning/gpu-utilization>.
- [52] SemiAnalysis. Training data and model storage requirements at scale, 2024. URL <https://semianalysis.com/>.
- [53] SemiAnalysis. AI chip supply constraints and multi-year agreements, 2024. URL <https://semianalysis.com/>.
- [54] Jaime Sevilla, Pablo Villalobos, Lennart Ho, et al. Training compute of frontier AI models grows by 4–5x per year. Epoch AI, 2024. URL <https://epoch.ai/blog/training-compute-of-frontier-ai-models-grows-by-4-5x-per-year>.
- [55] Claude E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423, 1948. URL <https://doi.org/10.1002/j.1538-7305.1948.tb01338.x>.
- [56] Brad Smith. The golden opportunity for American AI. Microsoft Official Blog, January 2025. URL <https://blogs.microsoft.com/blog/2025/01/03/the-golden-opportunity-for-american-ai/>.
- [57] Jakob Suchan, Mehul Bhatt, and Srikrishna Varadarajan. Commonsense visual sensemaking for autonomous driving. *Applied AI Letters*, 2(3), 2021. URL <https://doi.org/10.1002/ai12.56>.
- [58] The Information. Inside OpenAI’s infrastructure spending, 2024. URL <https://www.theinformation.com/>.
- [59] The Information. Microsoft’s NVIDIA GPU purchases, 2024. URL <https://www.theinformation.com/>.
- [60] The Information. OpenAI model deployment delays due to GPU shortages, 2024. URL <https://www.theinformation.com/>.
- [61] The Verge. OpenAI’s infrastructure commitments through 2035, 2025. URL <https://www.theverge.com/>.
- [62] Tom’s Hardware. Best GPUs for deep learning 2025, 2025. URL <https://www.tomshardware.com/reviews/best-gpus-for-deep-learning>.
- [63] TSMC. Chairman Mark Liu on CoWoS capacity constraints. TSMC Earnings Call, 2024. URL <https://investor.tsmc.com/>.
- [64] U.S. Department of Energy. AI and data center energy use: Challenges and opportunities. Technical report, U.S. Department of Energy, 2024. URL <https://www.energy.gov/policy/articles/ai-and-data-center-energy-use>.
- [65] Utility Dive. Data center power demands strain US electric grid, 2024. URL <https://www.utilitydive.com/>.
- [66] Vast.ai. Cloud GPU marketplace pricing, 2025. URL <https://vast.ai/>.
- [67] Vodworks. How much does AI cost: A C-level breakdown for 2025. Epoch AI analysis, 2025. URL <https://vodworks.com/blog/ai-development-cost>.
- [68] World Resources Institute. Powering the US data center boom: The challenge of forecasting electricity needs, 2024. URL <https://www.wri.org/insights/us-data-centers-electricity-demand>.
- [69] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-MNIST: A novel image dataset for benchmarking machine learning algorithms. *arXiv preprint*, arXiv:1708.07747, 2017. URL <https://arxiv.org/abs/1708.07747>.