

FROM MONTHS TO MINUTES: A NEW ARCHITECTURE FOR DOMAIN AI MODELS

The New Patent-Pending
Adaptive AI Approach that
Needs No Labels, No Training
Runs, and No Waiting



EXECUTIVE SUMMARY

Artificial Intelligence (AI) has made remarkable strides in recent years, yet organizations operating in specialized, high-variability domains consistently encounter a frustrating reality: general AI methods fall short when the stakes are highest. Rare instances (such as specific defects) go undetected. Model updates take months. Expert knowledge gets lost in translation.

Adaptive AI is a patent-pending machine learning approach designed by eSmart

Systems to overcome these limitations. It enables continuous learning from small data, rapid adaptation without retraining, and direct integration of evolving domain knowledge — creating AI systems that adapt and improve alongside the humans who use them. This whitepaper explains what makes Adaptive AI genuinely different, why few-shot learning alone is insufficient, and how Adaptive AI compares to traditional approaches in real-world conditions.



1 THE PROBLEM WITH TRADITIONAL AI IN COMPLEX DOMAINS

To appreciate what Adaptive AI offers, it helps to understand precisely where conventional approaches break down.

Large Data Requirements and Class Imbalance

Traditional deep learning demands large, balanced labeled datasets. In most industrial and domain-specific settings, visual inspection of power grid assets being a prime example, the most critical conditions are often the rarest. A model trained on thousands of examples of common components may have only a handful of examples of a rare but serious defect. The result is an uneven accuracy profile: strong performance on frequent cases, poor performance precisely where it is needed most.

Static Models That Cannot Keep Up

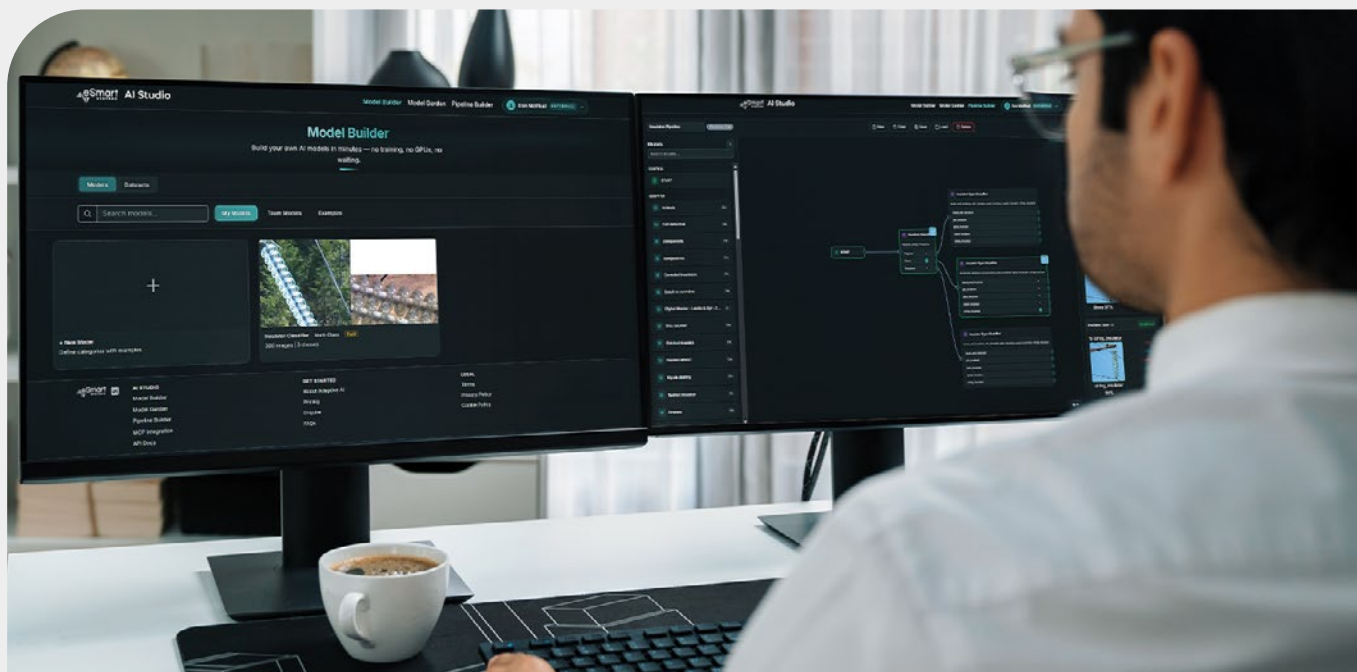
Once a conventional model is trained, its knowledge is effectively frozen. Improving it means collecting more data, relabeling, retraining, and redeploying; a cycle that can take weeks or months. In environments where definition of business rules (such as classification and severity) varies between customers, evolve over time, or must adapt to new regulatory standards, this rigidity is not a minor inconvenience. It is a fundamental incompatibility with how domain knowledge actually works.

Loss of Expert Intent

There is an inherent loss of meaning in the traditional AI pipeline. Domain experts articulate intent, datasets are constructed, labels are applied, and models learn statistical patterns from those labels. At each stage, nuance erodes. Edge cases are underrepresented. Subtle but important distinctions are diluted into aggregate statistics. The result is an AI that performs adequately on benchmarks but fails in the precise, ambiguous, decision-critical cases where accuracy matters most.

Slow and Uncertain Improvement Cycles

Retraining a traditional model is expensive, time-consuming and does not guarantee improvement. An iteration that takes weeks may yield minimal gains, fail to address the specific failure mode it targeted, or introduce regressions elsewhere. This uncertainty makes teams cautious about experimenting, which in turn slows the pace of improvement, the opposite of what dynamic domains require.



2

FEW-SHOT LEARNING: A VALUABLE INGREDIENT, BUT NOT A COMPLETE SOLUTION

Few-shot learning emerged as a promising response to data scarcity. By leveraging large pre-trained foundation models as feature extractors, few-shot methods allow a system to learn new concepts from only a handful of examples, an important advance over data-hungry conventional training.

But few-shot learning alone is not sufficient for real-world domain AI. Understanding why is essential to understanding what Adaptive AI adds.

Few-shot learning treats all examples as equally relevant

Most few-shot methods, such as basic prototypical networks, represent a class by averaging the feature vectors of all provided examples. This is a reasonable starting point, but it ignores the reality that some examples are far more informative than others. Averaging in irrelevant or misleading examples actively hurts performance.

Few-shot learning is typically a one-off process

You supply examples, a classifier is built, and learning stops. There is no built-in mechanism for continuous refinement, no loop that incorporates new feedback, and no natural way to handle evolving definitions.

The model does not get better in production, it simply is what it was when it was first configured.

Few-shot learning does not address the feature space

It assumes that the features produced by the underlying foundation model are suitable as-is. In practice, large foundation models produce high-dimensional feature vectors with significant redundancy. When combining features from multiple models, this problem compounds. Few-shot methods offer no principled way to select the relevant dimensions or combine multiple model outputs effectively.

Few-shot learning provides no guidance on what to learn next

Users are left to decide which additional examples to provide, without any system-level feedback about what would be most useful. This makes improvement slow, inconsistent, and dependent on user intuition rather than principled learning strategy.


In short, few-shot learning is one important building block. Adaptive AI is the complete architecture built around it.

Adaptive AI in Use – AI Studio by eSmart Systems

Launched in 2026 and built upon patent-pending Adaptive AI technology, AI Studio enables users to create new detection and classification models using just a few examples.

Example: Classification of Insulator Material on Powerlines

1, Attach Dataset -> 2. Define classes -> 3. Select Example Images of each -> 4. Save, Run & Test Model -> 5. Dynamically Update with Additional Examples



3

WHAT MAKES ADAPTIVE AI DIFFERENT

Adaptive AI combines few-shot learning with several additional mechanisms — each designed to address a specific limitation of conventional approaches. Together, they constitute not just a better algorithm, but a fundamentally different way of building and maintaining AI systems.

Foundation Models as Evolving Infrastructure

Adaptive AI uses foundation models; large pre-trained models such as vision transformers or CLIP (Contrastive Language-Image Pre-training by OpenAI)¹ to convert raw inputs into rich feature representations. Rather than treating these as fixed dependencies, Adaptive AI treats them as evolving infrastructure. When a newer, more capable foundation model becomes available, Adaptive AI automatically benefits without any retraining or relabeling. In a documented benchmark on crossarm material classification, simply switching from an older foundation model (CLIP, 2021) to a newer one (Meta's PE, 2025) boosted F1 accuracy from 0.72 to 0.89 in under one minute, with no additional data or engineering effort.

Dynamic Vector Partitioning

Large foundation models produce large, high-dimensional feature vectors. Using all dimensions indiscriminately invites the curse of dimensionality and buries the signal in noise. This problem becomes more acute when combining outputs from multiple foundation models, each with different architectures, training objectives, and feature distributions.

Adaptive AI addresses this with Dynamic Vector Partitioning (DVP), a patent pending² capability that identifies and prioritizes the most relevant parts of a model's representation for each task.

Rather than relying on a single model end-to-end, DVP enables Adaptive AI to intelligently combine strengths across models, making multimodal setups both practical and effective.

Adaptive Prototypes

Standard few-shot learning typically represents a class using a single, fixed summary of all examples. As more examples are added, this representation can become diluted, leading to diminishing returns or stalled performance.



¹ <https://openai.com/index/clip/>

² <https://www.esmartystems.com/news-and-resources/esmart-sy-stems-introduces-adaptive-ai-and-verify-ai-to-strengthen-grid-reliability-and-risk-mitigation/>




Adaptive AI replaces static representations with adaptive prototypes. Instead of treating all examples equally, the system emphasizes the most relevant evidence for each decision and adapts its representation accordingly. This enables better handling of variation within a class and allows concept quality to improve over time as new feedback is incorporated — rather than being averaged away.

The Continuous Adaptive Loop

Perhaps the most consequential difference between Adaptive AI and basic few-shot learning is what happens after the initial model is built.

Adaptive AI is designed as an ongoing dialogue between humans and the system. Users begin with a small set of examples, review results, and provide targeted feedback focusing on the system's most confident mistakes, its most uncertain correct predictions, and the edge cases near decision boundaries. Because concepts are updated by adjusting stored example representations rather than retraining model weights, each iteration is nearly instantaneous. Changes can be evaluated immediately, and unhelpful adjustments can be reversed just as quickly.

This low cost of failure changes the nature of AI development. Teams can experiment freely, iterate rapidly, and converge on high-performing models without the cautious, slow-moving cycles that traditional retraining demands.



GLASS INSULATOR

4

HEAD-TO-HEAD COMPARISON

Dimension	Traditional AI	Few-Shot Learning	Adaptive AI
Data requirements	Large-labeled datasets required; struggles with rare cases	Small initial samples can suffice, but relies on broad pre-training coverage	Learns from very few examples by design; handles rare and long-tail cases well
Adaptability	Static after training; requires full retraining to update	Adapts initially but is largely a one-off process	Continuously updatable; concepts refined in real time without retraining
Performance on rare cases	Poor. Data imbalance leads to uneven accuracy	Moderate. Depends heavily on pre-training domain coverage	Strong. Each concept is defined and refined independently
Iteration speed	Weeks to months per improvement cycle	Fast initial setup; slower for subsequent changes	Real-time; changes apply in seconds or minutes
Domain specificity	Requires large domain-specific datasets; hard to inject niche expertise	Constrained by what the pre-trained model already knows	Built for domain-specific intelligence; experts communicate intent directly through examples
Improvement over time	Each improvement risks degrading other classes	No built-in continuous improvement mechanism	Incremental, isolated improvements that compound over time

Table 1 - Performance Comparison of Model Types



5

PERFORMANCE IN PRACTICE

In a benchmark comparing crossarm material classification models, the results illustrate both the strengths and current position of Adaptive AI. Without making changes, Adaptive AI is surfing on the updates from the underlying

foundation model, still providing results well above traditional few-shot learning. This illustrates the strength of Adaptive AI and how it cements its position, utilizing improvements of underlying technology:


Model Type	Foundation Model	F1 Score ³	Development Time
Adaptive AI	CLIP (OpenAI, 2021)	0.72	< 1 minute
Adaptive AI	PE (Meta, 2025)	0.89	< 1 minute
Traditional AI	N/A	0.94	~1-3 days + annotation time

Table 2 - Model Type Performance

The traditional model achieves the highest F1 score, but at a cost that is orders of magnitude greater in time and effort. The Adaptive AI model, using the latest foundation model, closes the gap to within 5 percentage points while requiring no conventional training,

no labeled dataset construction, and no engineering iteration. Crucially, as foundation models continue to improve, the performance of Adaptive AI will continue to rise automatically. The traditional model will not.

³ The F1 score is a machine learning metric that measures a model's accuracy by calculating the harmonic mean of precision and recall, providing a balanced performance score between 0 (worst) and 1 (best). It is essential for imbalanced datasets, ensuring the model performs well without over-focusing on either false positives or false negatives.



6

STRATEGIC IMPLICATIONS

Adaptive AI is not simply a more efficient algorithm. It represents a structural shift in how organizations can develop and maintain AI capabilities.

Domain experts become direct contributors

Because Adaptive AI learns from examples and feedback rather than from abstract labels fed into opaque training pipelines, subject matter experts can build and refine models themselves. Specialized knowledge does not need to be mediated by data science teams or diluted through annotation processes.

Domain intelligence becomes a scalable, reusable asset


Over time, an Adaptive AI platform accumulates a growing library of discrete, honed concept detectors, each representing a specific condition, component, or failure mode. These concepts can be documented, reused across applications, and composed into larger AI systems and multi-agent pipelines as callable domain-specific tools.

Foundation models become a platform, not a ceiling

Rather than attempting to train one large model that encapsulates all domain knowledge, a slow, expensive, and often impractical endeavor, Adaptive AI treats foundation models as general-purpose infrastructure and adds domain-specific understanding as a modular, continuously evolving layer on top.

Time-to-value compresses dramatically

Where traditional AI projects measure their development cycles in months, Adaptive AI can produce deployable models in hours or days. More importantly, those models continue to improve in production rather than degrading as conditions change.



7

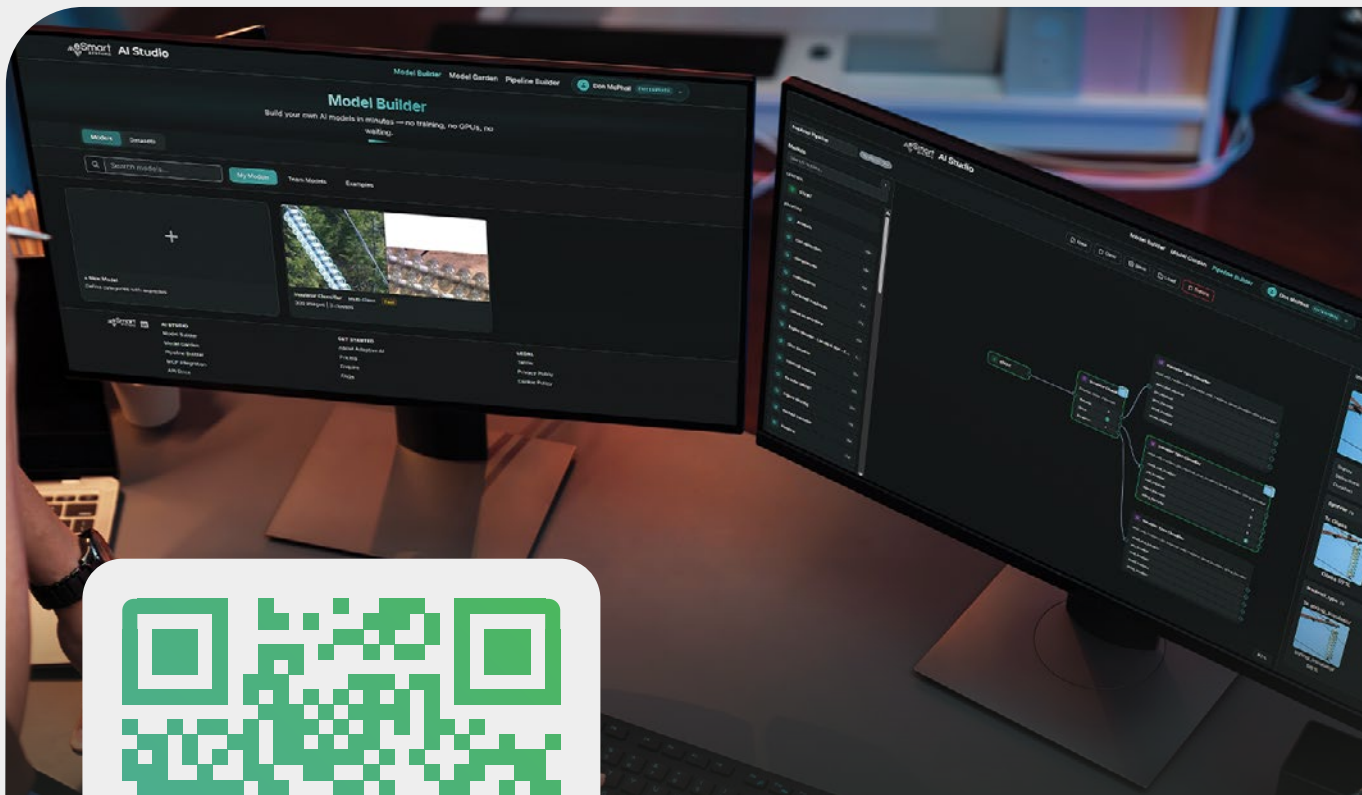
CONCLUSION

The limitations of traditional AI in complex, specialized domains are not incidental, they are structural. A paradigm that requires large datasets, tolerates only fixed class definitions, and improves only through costly retraining cycles is fundamentally mismatched with environments defined by variability, rarity, and evolving knowledge.

Few-shot learning addresses one dimension of this problem; data scarcity, but leaves the others unresolved. It provides no mechanism for continuous refinement, no principled approach to feature selection, and no guidance on how to make learning more efficient over time.

Adaptive AI addresses the full problem. By combining foundation model representations with Dynamic Vector Partitioning, adaptive prototypes, and a continuous human-in-the-loop feedback architecture, it delivers an AI system that learns quickly, adapts continuously, improves predictably, and stays aligned with the intent of the domain experts who use it.

In doing so, it transforms AI development from a slow, one-shot endeavor into an ongoing learning process; one that mirrors how human expertise actually grows and evolves.¹



¹ The results and technical descriptions in this whitepaper are based on Adaptive AI research and internal testing at eSmart Systems, contextualized with broader trends in foundation model adaptation and few-shot learning research.

To learn more about **Adaptive AI** and leverage it for your business, head to ai.esmartsystems.com