BlackIQ

PUBLIC WHITEPAPER / 2026

The Intelligence Layer

Architecture and economics of applied AI.

Data systems, model development, validation and controlled deployment in a common application framework.

BLACKIQ / NOVEMBER 2026

A technical and economic perspective on applied artificial intelligence.

PUBLIC PERSPECTIVE

From model capability to practical application.

The performance of an AI model and the reliability of an AI application are different questions.

Applied artificial intelligence combines statistical methods with data management, software interfaces, evaluation procedures and operational responsibilities. Performance measured during development is relevant, but it does not determine how a system will behave with new information or under different operating conditions. These distinctions are central to model development, quantitative research and enterprise software.

Data quality. Model selection. Validation. Operational control.

INTRODUCTION

An operating framework for applied AI

The commercial value of machine learning depends on how models are developed, tested and integrated into existing work.

Models can classify information, estimate future quantities and learn policies for sequential decisions. Their performance nevertheless depends on the observations available, the objective selected during training and the conditions in which a result is used. Errors in any of those areas can outweigh improvements in the underlying algorithm.

An application must therefore address problems that cannot be solved by model parameters alone. Data require consistent identities and timestamps. Experiments require meaningful comparisons. Results require documentation of their assumptions. Automated operations require defined permissions, limits and monitoring. These requirements affect cost, reliability and the ability to maintain software over time.

BlackIQ develops products around this relationship between research and application. Studio provides an environment for working with information and connected tools. Trade focuses on market research, model training, backtesting and the assessment of quantitative strategies. The products serve different purposes and cannot be reduced to a single learning method.

CENTRAL PROPOSITION

A model becomes operationally valuable when its inputs, performance, limitations and permitted uses can be managed as part of a complete process.

The following analysis begins with the distinction between model output and application performance. It then examines data management, research methods, model validation, reuse and the economic implications of integrating artificial intelligence into production workflows.

Intended readership

This publication addresses users, developers, researchers and organizations evaluating applied artificial intelligence.

READING MAP / CONTENTS

Research, application and operations

Eight chapters follow the development of an AI system from model output to operational and economic assessment.

01 From model output to decisions 5

The distinction between prediction, interpretation and application.

02 The intelligence layer 8

An architecture linking models, data and operating functions.

03 Data quality and context 11

Source identification, time consistency and data preparation.

04 Learning methods and selection 14

Supervised learning, reinforcement learning and world models.

05 Model evaluation and control 17

Independent testing, operational limits and authorization.

06 Reuse and interoperability 21

Versioning, interfaces, intellectual property and portability.

07 Economics of applied AI 24

Costs, adoption and the business case for integration.

08 Implementation and outlook 27

Technical priorities and conditions for wider adoption.

References and publication information follow the main chapters.

CHAPTER 01

From model output to decisions

A measured result is not, by itself, an operational decision.

Prediction, interpretation and use

01

01 / APPLICATION REQUIREMENTS

Model performance in practice

Evaluation scores describe behaviour under specified conditions; deployment introduces additional assumptions and constraints.

Consider a quantitative policy trained on historical market observations. A simulation can estimate how that policy would have performed under a given set of prices, costs and execution rules. It cannot establish that the same outcome will occur under future liquidity, volatility or market conditions. The result should be interpreted within the assumptions of the test.

The same distinction applies to information systems. A language model may produce a coherent report while omitting a material source or using an outdated document. Linguistic quality is not a substitute for factual accuracy, access permissions or suitability for the intended audience.

A production application includes interfaces between data, models, operating rules and users. These interfaces determine whether a numerical or textual output can be reviewed, reproduced and incorporated into a defined business process. The relevant unit of analysis is the process, not the model in isolation.

A DISTINCTION IN PRACTICE

Data define the observations. **Models** estimate relationships or select actions. **Tests** assess performance under stated assumptions. **Permissions** determine what an application may do. **Operations** determine how the system behaves over time.

Integration costs and information loss

Organizations often train a model in one environment, evaluate it in another and deploy it through a separate software system. Each transfer creates the possibility of inconsistency: a changed input definition, a new model version, a different fee assumption or an authorization setting that does not correspond to the original experiment.

Maintaining a consistent record of inputs, configurations and results reduces the cost of repeated work. It also makes differences between experiments easier to identify and operational failures easier to investigate.

Model quality is not operational readiness.

Application design principle

Decision authority remains explicit

An application may recommend, classify or execute, depending on its design and granted permissions. In each case, the decision criteria, acceptable risks and authority to act must be defined outside the model's optimization objective. Numerical performance cannot establish the right to perform a transaction or access restricted information.

Three stages of application

A model produces a measurement or recommendation. An analyst interprets that result in relation to the available information and uncertainties. A person or an authorized system then decides whether and how to proceed. These stages may be closely integrated in software, but they have distinct responsibilities.

01 / FROM MODEL OUTPUT TO A DECISION

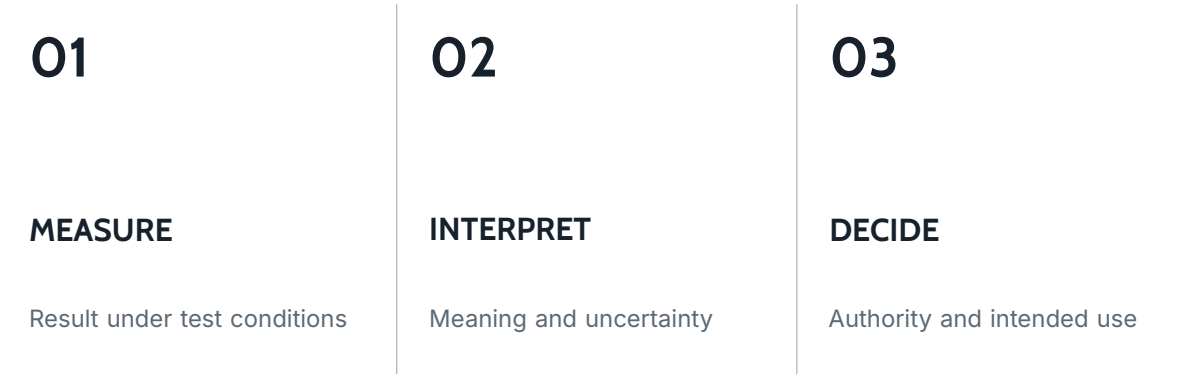

Model output, analytical interpretation and authorization serve separate purposes.

That distinction is particularly important in applications involving financial exposure or sensitive information. Model accuracy cannot determine whether a proposed action is permitted or appropriate.

The intelligence layer

A common application framework connects specialized AI capabilities.

Data, methods, validation and operations

02

02 / APPLICATION ARCHITECTURE

Connecting distinct technical functions

An intelligence layer organizes the interfaces between information, learning methods, application software and operating rules.

The term *intelligence layer* is used here for an application architecture rather than a new model class or network protocol. Its purpose is to connect established capabilities without assuming that their underlying algorithms are interchangeable. Data services, model runtimes, retrieval, policy controls and user interfaces address different technical requirements.

A language model can generate text or code. Retrieval systems can supply additional documents at inference time. Reinforcement-learning methods can optimize actions over a sequence of observations. Conventional application rules can impose access restrictions and operating limits. Combining these components requires reliable interfaces and a consistent understanding of the task.

02 / FUNCTIONS OF AN AI APPLICATION

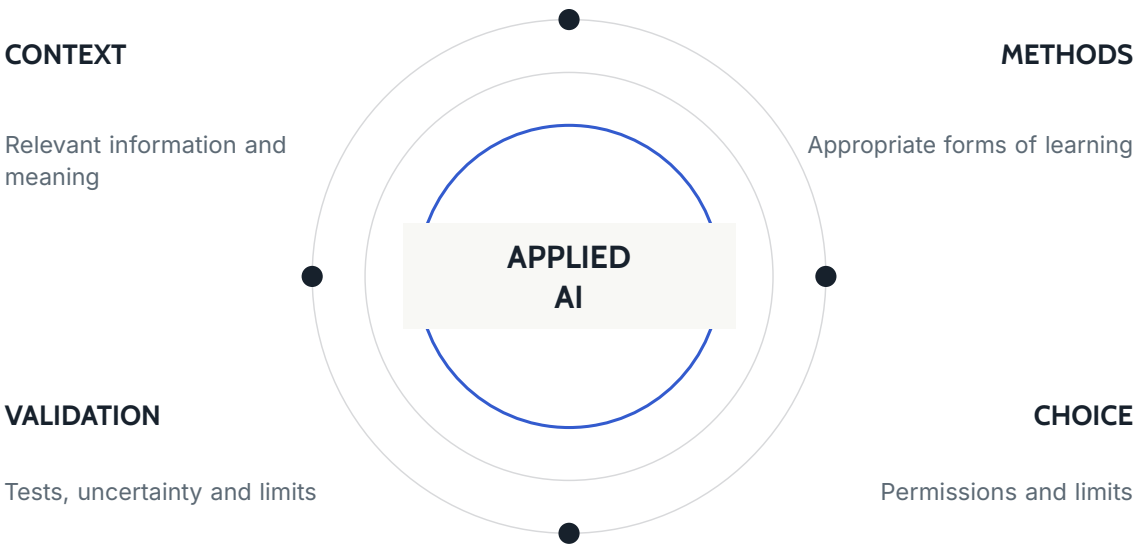

Application requirements considered together: context, methods, validation and operating decisions.

Two application areas

BlackIQ Studio focuses on information work and connected software tools. The relevant design questions include document access, information retrieval, source relevance, user permissions and the relationship between generated output and underlying records. Retrieval-augmented generation is one established research approach for combining model generation with external information [2].

BlackIQ Trade addresses quantitative market research. It brings market observations, training configurations, model comparisons and simulated performance into a common research process. A trading strategy requires separate consideration of data availability, transaction costs, market structure and the conditions governing any authorized execution.

Method selection depends on the application

Generating a document, retrieving an internal policy and estimating a trading action involve different objectives. They may differ in their observation spaces, time horizons, error costs and regulatory implications. An integrated platform can provide consistent controls and presentation while retaining appropriate methods for each domain.

This approach follows established scientific work rather than claiming a universal algorithm. Transformer architectures support sequence representation [1]. Retrieval methods connect generation with external information [2]. Reinforcement learning addresses decisions with delayed consequences [3]. The engineering task is to choose and integrate these methods where their assumptions fit the intended application.

Data quality and context

Data integrity determines the validity of downstream analysis.

Identity, timing and transformation

03

03 / DATA ENGINEERING

Defining reliable source data

A consistent dataset identifies what was observed, when it was observed and how each record was transformed.

Market data demonstrate the importance of precise definitions. A price series depends on the trading venue, instrument, unit of account, interval and rules for missing or duplicate records. On-chain observations require additional interpretation of transactions, event logs, pools, token precision and transaction finality. Two series with identical labels are not necessarily economically comparable.

The relevant question is whether an observation can be connected to the input supplied to an experiment. Data preparation may involve normalization, filtering, interval aggregation and corrections to historical records. Those operations can change the statistical properties of a series and should be documented.

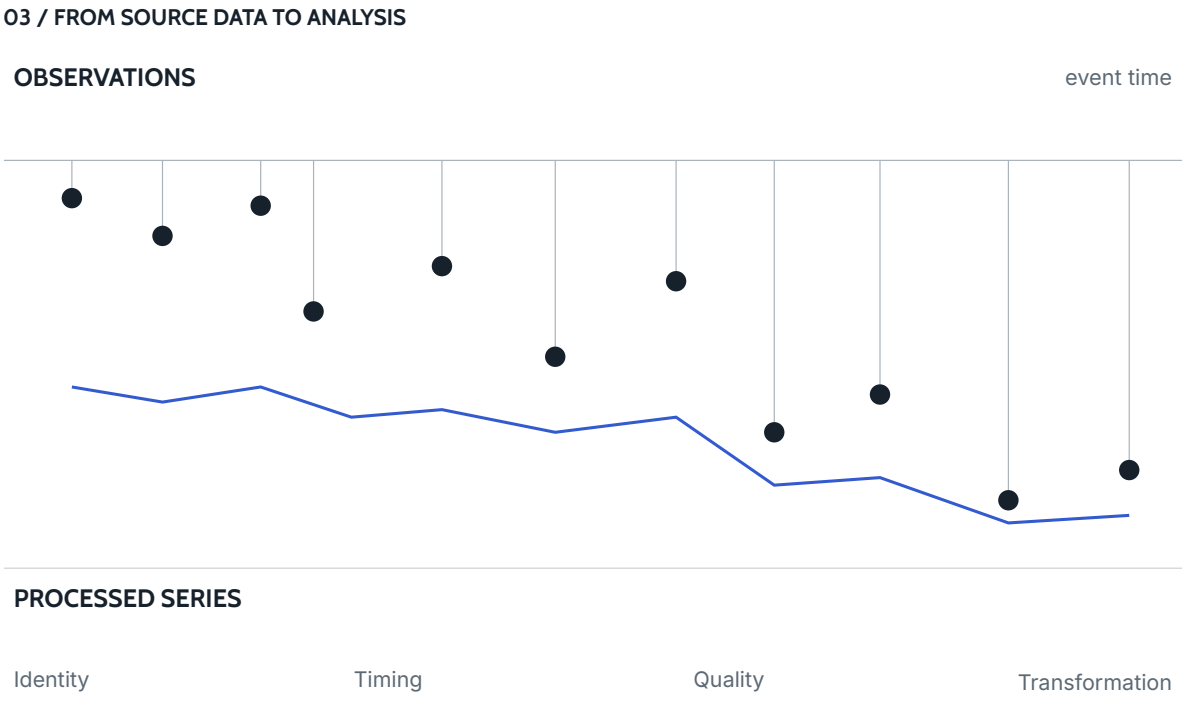

Schematic observation sequence and transformed series. Values are illustrative, not historical prices.

The effect of data errors

An incorrect base/quote convention, misplaced timestamp, stale price or incomplete interval can alter calculated returns and model inputs. Such errors may appear as model instability

or misleading simulated performance rather than as an explicit data failure. Data validation should precede statistical analysis and remain available when results are reviewed.

External datasets

Research often requires information that is not supplied by a platform. Support for user-provided datasets introduces requirements for numeric types, units, timestamps, missing observations and schema compatibility. A valid import process should distinguish the original dataset from its normalized form and from features generated for a particular experiment.

Versioned input definitions enable reproducible comparisons. Where the data source or transformation changes, the difference must be identifiable rather than silently absorbed into an existing experiment.

Local and hosted execution

Supported training workloads may run locally or in hosted computing environments. The relevant differences concern hardware, compute capacity, data transfer, storage and the services used during execution. A shared configuration interface does not imply identical runtime characteristics.

Local execution can be appropriate when the device and model support it. It does not, by itself, establish that a workflow is private, secure or independent of network services. Those properties depend on storage, synchronization and the concrete software implementation.

DATA MANAGEMENT

The original input, normalized dataset, experiment configuration and trained artifact should remain distinguishable. That separation makes errors easier to locate and experiments easier to reproduce.

Learning methods and selection

Different objectives require different statistical and computational methods.

Supervised, sequential and model-based learning

04

04 / RESEARCH METHODS

Method selection

Model choice should reflect the prediction target, time horizon, available observations and costs of incorrect decisions.

Supervised learning estimates relationships from labelled or otherwise defined training examples. Representation learning can identify features that are not manually specified. Transformer-based architectures have extended sequence modeling through attention mechanisms [1]. Their general applicability does not remove the requirement to evaluate performance under the data distribution relevant to a specific use.

Sequential decisions require a different formulation. An action may change subsequent observations and therefore the opportunities available later. Reinforcement learning studies this interaction through states, actions, rewards and policies [3]. Its relevance depends on the quality of the state representation and the degree to which the environment models the intended operating conditions.

04 / RESEARCH METHODS BY OBJECTIVE

01	ESTIMATE	Statistical relationships learned from observed examples.	
02	ACT	Sequential policy optimization; SAC is one example.	
03	SIMULATE	Learned environmental dynamics; Dreamer is one example.	
04	RETRIEVE	External documents, memory and software interfaces.	

Families of published research methods; this is not a list of product deployments.

Soft Actor-Critic

Soft Actor-Critic (SAC) is an off-policy reinforcement-learning method for continuous-action problems. Its maximum-entropy objective combines expected returns with an entropy component that influences policy exploration [4]. SAC provides a useful research basis for selected sequential control problems; it does not independently establish the profitability or robustness of a market strategy.

A suitable method must be assessed against available data, action constraints, training stability, computational requirements and the cost of an incorrect outcome. A more sophisticated algorithm is not necessarily more effective than a simpler baseline.

World models and Dreamer

Model-based reinforcement learning attempts to represent aspects of an environment's dynamics. A learned model can be used to simulate possible transitions and support policy training. The Dreamer research family provides a published example of this approach; DreamerV3 demonstrated results across diverse control tasks [14].

Those findings are contributions to the research literature. They do not demonstrate that learned market dynamics are sufficiently stable for trading, nor do they establish which algorithms are deployed in a particular product environment. Teacher-guided learning and policy distillation represent other research techniques whose suitability must be established through separate testing.

Objective design and market conditions

An optimization objective is a simplified expression of the desired outcome. In market research, the difference between simulated reward and realizable performance can arise from trading fees, slippage, liquidity, latency and incomplete information. Optimization can amplify a poorly specified objective rather than correct it.

Experiment comparison should therefore include appropriate baselines, cost assumptions and independent assessment. The selection criterion should be stated before choosing the model presented as the strongest candidate.

Model evaluation and control

Independent testing and authorization govern the transition to operation.

Validation, limits and oversight

05

05 / VALIDATION AND AUTHORIZATION

Interpreting simulated performance

Historical performance is conditional on the test period, model selection procedure and assumptions about execution.

A backtest evaluates a strategy against a specified historical environment. Results can be distorted by data leakage, repeated parameter searches, inappropriate sampling or assumptions that exclude material costs. A nominal holdout is not independent if the same observations influenced the final choice of model or parameters.

Studies of deep reinforcement learning report sensitivity to random seeds and experimental settings [5]. A technically credible comparison therefore identifies the sampling method, training period, baselines and variability across runs. The evaluation protocol is part of the research result, not a secondary reporting detail.

05 / CHRONOLOGICAL MODEL ASSESSMENT

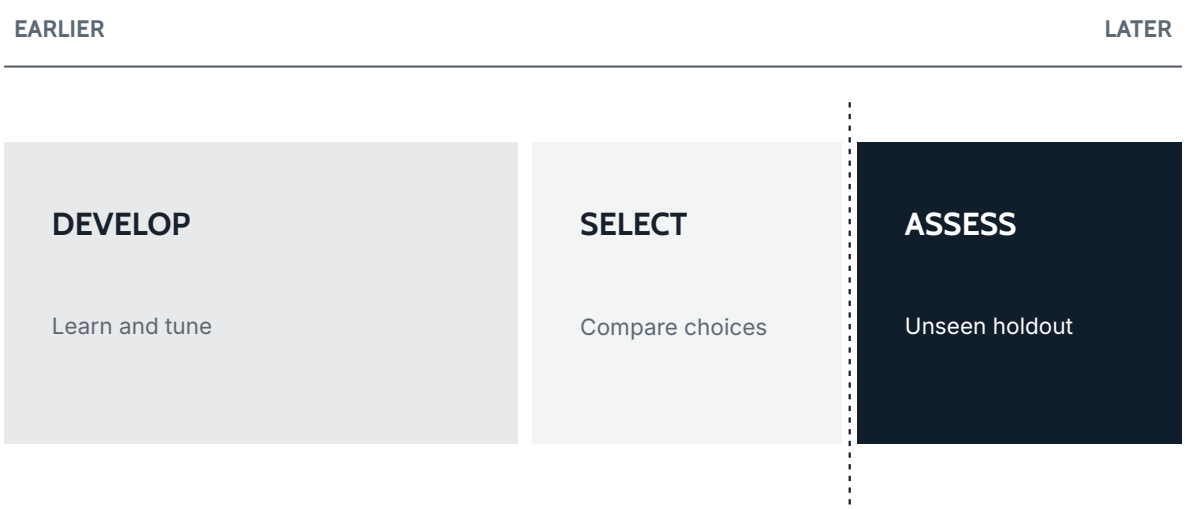

Freeze design before inspecting the final holdout.

Temporal separation of development, selection and final testing. The schematic does not report performance.

Independent assessment

A well-defined protocol separates development, model selection and a final assessment on previously unused observations. For time-dependent data, the temporal order of available information must be respected. Walk-forward experiments and comparisons across different conditions can provide a more demanding assessment than a single

favorable test interval.

Transaction costs, slippage, order size, execution delay and market liquidity should reflect the intended application. A model that performs well under one set of assumptions may perform differently when those assumptions change.

INTERPRETATION OF RESULTS

A positive simulation result establishes performance under the assumptions of that simulation. It does not establish future returns, immunity to changing market conditions or suitability for every portfolio.

Authorization is a separate decision

A model may recommend a course of action without being authorized to carry it out. Where execution is supported, the application must distinguish research output from a transaction request. Asset eligibility, budget limits, duration, user permissions and the ability to revoke access are operational requirements independent of model performance.

Access controls should apply to the transaction that is actually submitted, not only to the interface that presents the recommendation. Missing market data, failed services and changes in authorization require defined operating behavior. Features and restrictions differ between supported environments.

06 / VALIDATION AND AUTHORITY

01	02
VALIDATION	**AUTHORITY**
Was the model tested under conditions that matter?	Who decides whether an action may happen?

THE TWO CONDITIONS BELONG TOGETHER Review / limit / stop

Validation establishes performance under stated conditions; authorization determines permissible actions.

Reporting and accountability

An experiment record should identify the model and dataset versions, evaluation interval, benchmark, transaction-cost assumptions and relevant failure conditions. Reports should describe both central results and variability rather than presenting only the best observed run.

Standards such as the NIST AI Risk Management Framework provide a general reference for assessing the organization and management of AI-related risks [6]. They do not, by themselves, establish compliance or validate the performance of a particular product.

Human review and operating responsibilities

The appropriate division of responsibility depends on the application. A research system may present alternatives for review. A service permitted to act requires explicit authorization, accountable operators and a mechanism to suspend activity. The interface should disclose the scope of the permission and the information used to reach a recommendation.

Reuse and interoperability

Models require documented interfaces, versions and rights to support reuse.

Beyond a single training run

06 / MODEL LIFECYCLE

A reusable model is a maintained software artifact

Model parameters alone do not define the inputs, operating assumptions or permitted applications of a trained system.

A model file can preserve learned parameters without specifying how the input data were prepared, which version of a runtime was used or how performance was assessed. Reuse requires a stable identity for the artifact and a description of its input and output expectations, dependencies, tested conditions and rights of use.

Software adapters and extension interfaces make it possible to connect additional data sources and specialized methods without rebuilding the entire application. Their value depends on version compatibility, clear interfaces and a defined responsibility for maintenance.

Models, datasets and experiments

For a researcher, reuse may mean returning to a previous run or comparing candidates under consistent assumptions. For a developer, it may mean integrating a source through a supported adapter. For a team, it may mean exchanging an experiment definition that another researcher can inspect or repeat.

These activities require more than transferring an undifferentiated file. Input definitions, configuration changes and the test conditions attached to a result should remain identifiable.

Rights and future distribution

Sharing or licensing specialized models raises questions about ownership, permissions, suitability, pricing and responsibility for updates. A marketplace for models or extensions is part of BlackIQ’s longer-term direction, not a statement that every form of model distribution or rental is currently available.

DISTRIBUTION REQUIREMENTS

Model distribution depends on versioned artifacts, defined interfaces, clear usage rights and adequate documentation of assessed conditions. Commercial terms cannot substitute for technical suitability.

Compatibility is not certification

A documented model can still be inappropriate for a different dataset, market or operating environment. Compatibility with a software interface does not guarantee statistical performance or regulatory suitability. Such judgments remain dependent on the intended application and its constraints.

Distribution mechanisms must also distinguish ownership of a model from permission to act on behalf of another user. These are separate legal and technical relationships.

07 / MODEL REUSE REQUIREMENTS

01	**ORIGIN**	The data and assumptions behind it
02	**METHOD**	What was trained and how
03	**VALIDATION**	Test design and reported results
04	**BOUNDARIES**	Version, rights and permissions

A reusable model requires source information, method definitions, test results and usage conditions.

Consistent execution interfaces

A shared experiment description can support both local and hosted computing when the relevant runtimes are available. Hardware support, storage and network behavior still differ. Versioned configurations make those differences easier to document.

Economics of applied AI

The business case depends on productivity across an entire process.

Costs, adoption and investment

07

07 / ECONOMIC ASSESSMENT

The full cost of an AI application

Inference cost is only one component of the resources required to deploy and maintain an AI application.

The direct cost of a model call is relatively easy to observe. Data acquisition, preparation, application integration, quality review, corrections and maintenance are less visible but can determine whether the investment is justified. A low-cost model can be expensive to use when its outputs require substantial downstream correction.

Economic studies of information technology emphasize complementary investments in organizational practices, skills and workflow redesign [9]. Research on AI adoption makes a similar distinction between access to models and measurable changes in productivity [7,8]. As a result, the relevant comparison is between alternative ways of completing a process, rather than inference prices alone.

08 / OPERATING COST CATEGORIES

01	DATA	Preparation and access
02	COMPUTE	Training and inference
03	REVIEW	Verification and correction
04	INTEGRATION	Deployment and upkeep

Categories required for a process-level cost assessment. The diagram does not indicate relative expenditure.

Costs should be assessed across development and operations: data preparation, model training or inference, software integration, review and continuing maintenance. The appropriate measures depend on the application, but can include time, error rates, compute use and the resources required to correct failures.

Adoption statistics and their limits

The OECD reported AI use among firms in countries with available observations of 8.7% in 2023, 14.2% in 2024 and 20.2% in 2025 [11]. These estimates concern reported adoption under the OECD's statistical coverage. They do not establish productivity gains, revenue growth or performance for BlackIQ users.

09 / REPORTED ADOPTION BY FIRMS

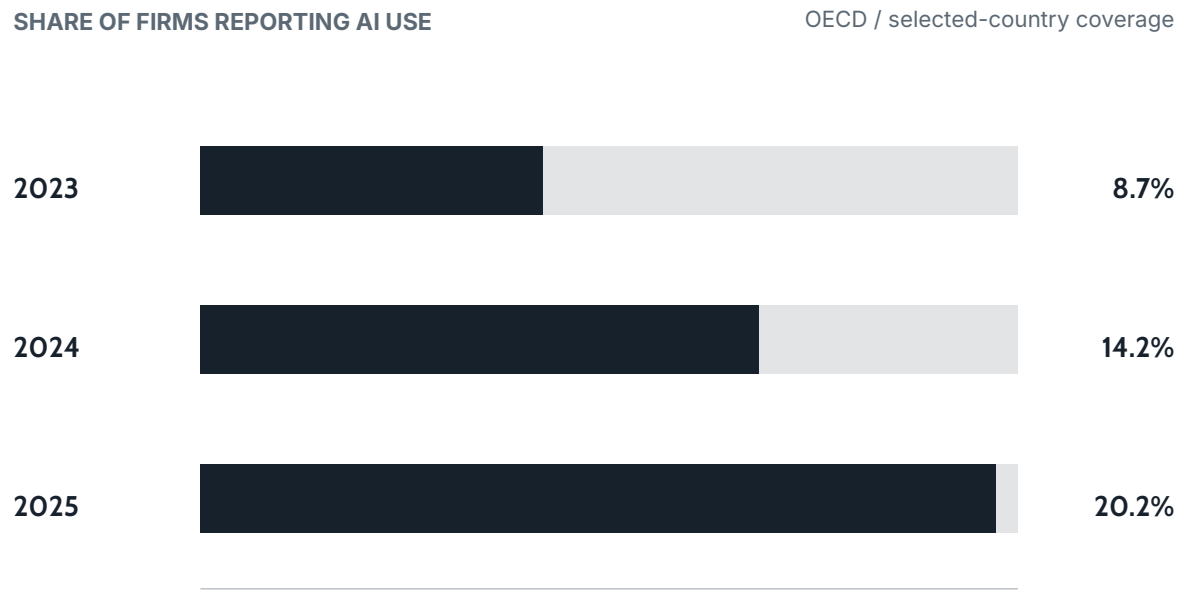

OECD data for countries with available firm statistics [11]. Adoption rates do not measure productivity or platform results.

Potential sources of economic value

A common application framework may reduce repeated data preparation, fragmented software integration and the cost of comparing or reproducing model experiments. For individual users, the benefit may be time saved during research. For developers, it may be a reusable interface. For organizations, it may be more consistent documentation and process management.

These outcomes are propositions to be assessed through actual use. Their value depends on implementation costs, the alternatives available and whether an improvement persists outside a demonstration environment. The economics of platform distribution also depend on the rights and incentives of different participants [12,13].

Implementation and outlook

Further adoption requires technical reliability and credible economic results.

Research, product development and policy considerations

08

08 / OUTLOOK

Requirements for wider application

The development of applied AI will be measured by reproducible performance, reliable interfaces and economic results in defined settings.

Improving model capabilities remains important. Production applications also require work in data engineering, evaluation, software maintenance and operational controls. These disciplines influence whether a new model can be introduced without changing the meaning of existing results or increasing risk in ways that are difficult to observe.

Information-oriented workflows and quantitative market research demonstrate different parts of the same engineering challenge. Studio places greater emphasis on access to information and connected tools. Trade places greater emphasis on temporal datasets, experimental comparison and the separation of simulation from authorized execution. Both require domain-specific methods rather than a single generic approach.

Research in reinforcement learning, world models and model distillation expands the available methods. Adoption should follow demonstrated suitability for a particular task. A published research result does not establish a production feature, and a promising simulation does not establish future financial performance.

Measure performance in the intended application.

Conclusion

ASSESSMENT CRITERIA

Before adopting an AI application, establish the origin of its inputs, the conditions used to assess it, the responsibilities assigned to its operators and the costs of maintaining it. These questions provide a practical basis for procurement, development and public discussion.

REFERENCES

Scientific and economic literature

Selected publications on model architectures, reinforcement learning, evaluation and the economics of technological adoption.

[1] Vaswani, A. et al. (2017). *Attention Is All You Need*. NeurIPS. Source.

[2] Lewis, P. et al. (2020). *Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks*. arXiv:2005.11401.

[3] Sutton, R. and Barto, A. (2018). *Reinforcement Learning: An Introduction*, 2nd ed. MIT Press. Publisher.

[4] Haarnoja, T. et al. (2018). *Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor*. ICML, PMLR 80. Paper.

[5] Henderson, P. et al. (2018). *Deep Reinforcement Learning That Matters*. AAAI. DOI.

[6] NIST (2023). *AI Risk Management Framework 1.0*. AI 100-1. DOI.

[7] OECD, BCG and INSEAD (2025). OECD research publication on the adoption of artificial intelligence in firms. DOI.

[8] OECD (2025). *The Effects of Generative AI on Productivity, Innovation and Entrepreneurship*. DOI.

[9] Brynjolfsson, E., Hitt, L. and Yang, S. (2002). *Intangible Assets: Computers and Organizational Capital*. Study.

[10] OECD (2024). *The Impact of Artificial Intelligence on Productivity, Distribution and Growth*. DOI.

[11] OECD (2026). *AI use by individuals surges across the OECD as adoption by firms continues to expand*. Jan. 2026. Announcement.

[12] Rochet, J.-C. and Tirole, J. (2003). *Platform Competition in Two-Sided Markets*. JEEA 1(4). DOI.

[13] Comunale, M. and Manera, A. (2024). *The Economic Impacts and the Regulation of AI: A Review*. IMF WP 2024/065. DOI.

[14] Hafner, D., Pasukonis, J., Ba, J. and Lillicrap, T. (2025). *Mastering diverse control tasks through world models*. Nature 640, 647–653. DOI.

PRODUCT INFORMATION

Current product descriptions and regional availability are published at blackiq.ai. Supported features and permissions depend on the selected environment and applicable terms.

PUBLICATION INFORMATION

Scope and limitations

This publication provides an overview of applied AI methods and the application framework pursued by BlackIQ. Scientific references describe published research, while product descriptions summarize public information. No original controlled benchmark of BlackIQ's performance is presented here.

Availability depends on the software version, deployment environment, supported data and user permissions. Research on learning algorithms must not be interpreted as confirmation that a particular algorithm is deployed in every product or supported market.

This publication is not investment advice, an offer of securities or a guarantee of trading performance. Model-based strategies involve financial risk. References to general standards do not constitute a claim of regulatory approval or security certification.

Source categories

Published research, architectural concepts, current product descriptions and longer-term development plans serve different purposes. A cited study supports a scientific statement within its scope; it does not certify the implementation of a commercial product. Proposed features should not be interpreted as available services unless explicitly confirmed in current product information.

Publication

BlackIQ Public Whitepaper, November 2026. Prepared as a public technical and economic overview for users, researchers and other interested readers.



BLACKIQ / THE INTELLIGENCE LAYER

Public Whitepaper, November 2026

9 November 2026 | blackiq.ai