

The cost of creating an AI proof of concept

Intro

Today, we are witnessing AI-powered software that senses, reasons, acts, and adapts to new information. It starts by 'learning' from large sets of diverse data training models to recognise patterns and insights. Once the model has been trained, it can then apply what it has learned to new or unseen data, such as converting spoken words into text and sound, spotting anomalies in images, or detecting unusual patterns.

Although AI algorithms have been around for quite some time, recent advancements in AI technology have snowballed within various industries. This rapid growth is mainly due to significant drops in processing and data storage costs. At the same time, improvements in AI algorithm design, particularly in neural networks, have led to more accurate model training.

Such developments accelerate AI adoption, opening up many opportunities for organisations to make more informed decisions and improve processes, resulting in tangible business outcomes.

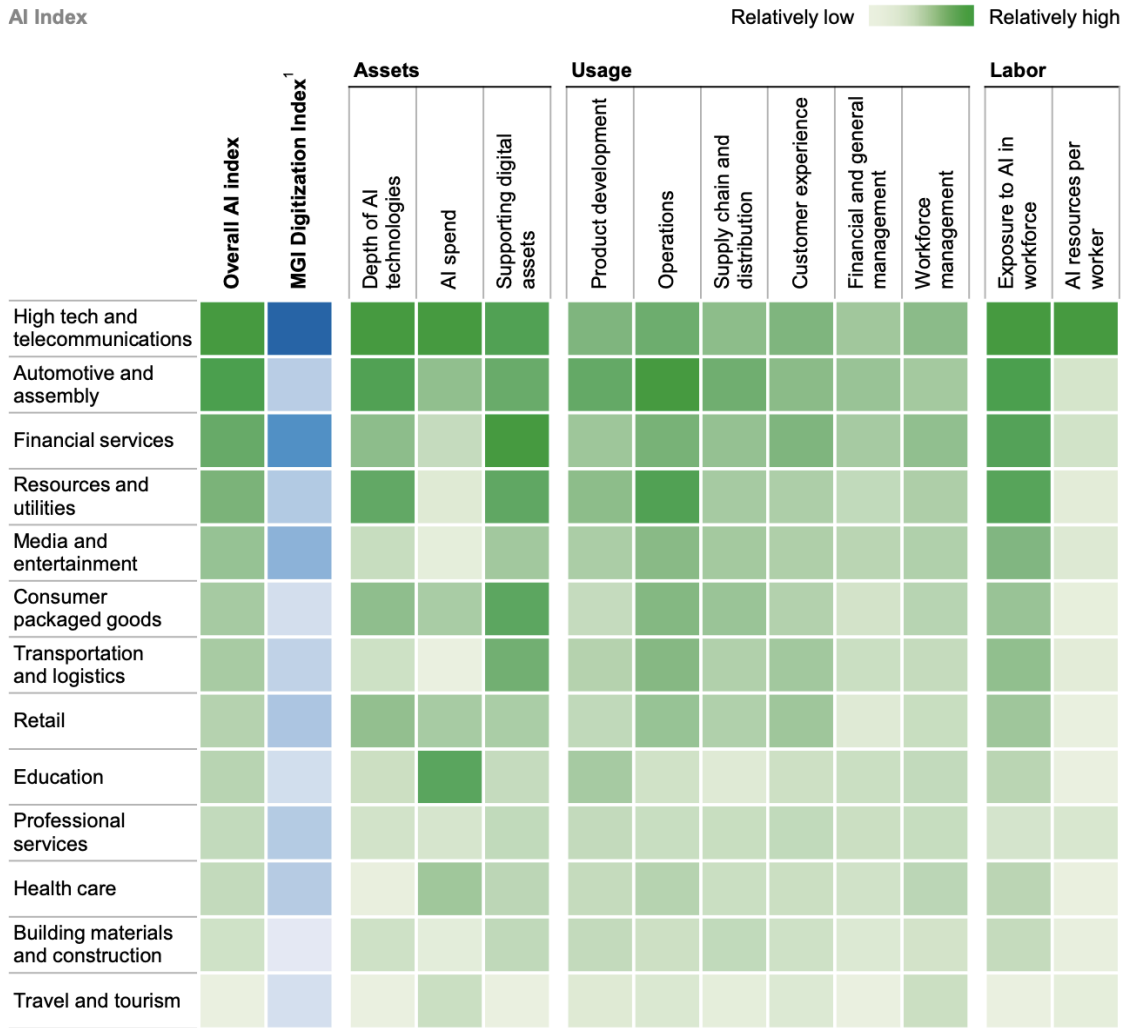
For many organisations, developing an AI Proof of Concept is the first step toward integrating AI into their operations. However, the costs associated with creating a PoC vary greatly depending on several factors, including project scope, data requirements, and the level of AI sophistication.

What is a proof of concept?

A proof of concept or PoC is a working solution, though limited in scope, that is created to test and evaluate whether an idea or approach is feasible under real-world conditions. It follows specific criteria, helping to confirm that requirements are understood and success is achievable.

In the context of an AI project, PoCs allow decision-makers to:

- < Deliver quick, tangible results
- < Gain practical experience and knowledge
- < Assess hardware, software, and service options
- < Identify and address potential data flow issues
- < Highlight how AI might impact IT infrastructure and the broader organisation
- < Build confidence in AI by demonstrating its value



The heat map visually compares the AI and digitisation levels within the industries

High-tech and telecommunications, financial services, and automotive industries lead the way with strong AI integration in areas like product development, customer experience, and workforce management. These industries have invested heavily in AI technologies and digital assets, showing high exposure of their workforce to AI and having abundant AI resources per worker.

However, this level of success didn't happen overnight—it was built on carefully planned Proof of Concept developments. AI PoC development is vital for businesses to test, validate, and refine AI applications before committing to full-scale implementation.

The phases of AI PoC development

Step 1. Clarify your AI opportunities

Before jumping into AI projects, you must clearly define what you aim to achieve, why it matters to your business, and how you'll know if it's successful. If you haven't yet pinpointed the key areas where AI can be helpful, start by assessing where it could bring the most immediate value:

- < Look at how others in your industry are using AI
- < Identify any pressing issues in your business that AI could help solve or areas where it might bring real improvements
- < Make use of the expertise you already have within your team

Once you have a list of potential areas where AI could make an impact, evaluate each one using a few important criteria. This doesn't have to be a lengthy process, but it's helpful to ask these questions to ensure you've planned well and aren't rushing into an AI project:

- < Have you clearly identified the problem you want to solve and how you will measure success? Have you already explored other solutions and ruled them out in favour of AI?
- < Is the scope of the project manageable? For example, can you outline the data needed, key components, and who will be affected? Will this be part of a larger solution?
- < Do you have the necessary technology, resources, and funding? Can you access the data you need without any roadblocks—whether technical, contractual, or otherwise?
- < Is the expected impact of the project significant enough to justify the effort? It's important to focus on visible wins that will build trust in AI and engage stakeholders.
- < Are the timelines realistic? Is there a clear, motivated delivery team with the time and skills needed to carry the project through?
- < Does the project fit into a broader AI or data strategy for your company? What existing infrastructure can support it?
- < Finally, if the Proof of Concept is successful, is there a plan for what comes next? Do you have funding to scale or maintain the solution? Is your IT team prepared to support it?

These questions help you weigh a solution's value, cost, and risk, which is the foundation of a strong business case—though for simpler PoCs, a formal business case might not be necessary.

Step 2. Define the problem and understand your data

Once you've pinpointed and assessed your AI opportunity, it's time to understand the unique problem you're going to solve. Start by categorising the issue into broader areas such as reasoning, perception, or computer vision.

For many organisations, especially those just starting with AI, a great challenge is having the right skills available internally.





This stage of your AI process is also a good moment to consider some technical questions that might affect your solution:

- ◀ Do you have specific hardware or software preferences? What are the reasons behind these choices (like benchmark data, total cost of ownership, or preferred suppliers)?
- ◀ Do any security, regulatory, or data considerations lead you to prefer on-premise systems over cloud solutions?
- ◀ Will your solution be locally managed, or will it be hosted in a data centre?
- ◀ How frequently will you receive new data for training and inference, and in what quantities?
- ◀ What measures will you take to ensure the security of raw data and insights while stored and during transmission?



Step 3. Design and implement the solution

The next step is designing and implementing the solution you tested in the Proof of Concept. AI solution architecture may be presented as a technology stack:

COMPONENT	DESCRIPTION
 Foundational products and systems Infrastructure	The core hardware and software systems that support AI operations.
 AI-specific software	Software solutions tailored to operate and manage the AI infrastructure.
 Frameworks for AI	Development frameworks that enable AI capabilities and support the planned solution.
 Visualisation and front-end software/hardware	Tools for visualising data and interacting with the AI solution.

At this point, you might be considering whether to build, buy, or repurpose existing hardware and software and use cloud services.

Even if the infrastructure and software are developed according to best practices, consider the specific requirements of AI. One critical aspect is the need for a continuous supply of high-quality data. Data scientists should collaborate with IT systems architects to design the deployment architecture, ensuring everything from the data centre to edge devices is accounted for - software integration, network connectivity, and physical considerations. We recommend testing multiple options using a test-and-learn approach to maximise learning opportunities.

Once the foundational elements are in place, you can focus on other AI-specific components of the solution, including model building, training, and tuning.

Model building

Building models is at the heart of AI work. Data scientists use training data and manage parameters to run iterative tests. This process helps them check for initial accuracy before moving on to broader training and fine-tuning.




Training and tuning

Training and tuning are the most demanding processes of the AI workflow. During this phase, data scientists determine the most efficient parameters for their models based on the available training data while addressing traditional IT issues like job scheduling and infrastructure management.





Step 4. Assess the business value

As you move forward, you'll want to evaluate your Proof of Concept against the criteria you defined during the solution design phase. For engineers, these criteria should be clear and measurable, allowing for continuous testing, ideally through automated processes.

Here is some key evaluation criteria focused on business value:

EVALUATION CRITERIA	DESCRIPTION
 Accuracy	Assess if the solution provides correct results and insights and whether these results can be consistently repeated.
 Completeness	Evaluate if the solution effectively employs all available data sources.
 Timeliness	Determine if insights are delivered when they are needed.

In addition to assessing business value, it's important to evaluate how well the solution functions overall:

EVALUATION CRITERIA	DESCRIPTION
 Scalability	Will the solution still perform effectively as data volumes or user numbers increase, whether gradually or suddenly?
 Compatibility	Can the solution easily integrate with third-party data sources and services using standard protocols?
 Flexibility	Can the solution adjust to new requirements if data needs or models change?
 Engineering	How easy is it to troubleshoot and correct any inaccuracies in the outputs generated by a trained model?

You should also consider explainability, which relates to the quality of decision-making. The criteria for explainability include:

- **Bias:** How can you ensure the AI system doesn't reflect biases present in the training data, the model itself, or even the objectives defined by its human creators?
- **Fairness:** If decisions are made using the AI system, how can you confirm that these decisions are fair? What does "fair" mean in this context, and for whom?
- **Causality:** Can the model make correct inferences and explain the reasoning behind these conclusions?
- **Transparency:** Does the AI provide the insights presented in a way that users can easily understand? How can users challenge or question the findings?
- **Safety:** How can users trust the reliability of the AI system, even if they don't fully understand how it arrives at its conclusions?

Step 5. Expand the proof of concept

Now that you've clearly defined the problem, designed the solution, and set up the data, it's time to look at scaling up your proof of concept after building and testing it. Here's what to consider moving forward.

Positive user experiences can create a growing demand for your solution, which is a great sign. However, if too many people show interest too quickly, it can overwhelm the PoC. To keep it running smoothly and to support a wider AI strategy, consider the following steps.

Five elements of successful AI transformations



Use cases/
sources of value




Data
ecosystems



Techniques
and tools



Workflow
integration



Open culture
and organization

- **Enhance inference:** Scaling up AI doesn't always follow a straight line. For example, if you upgrade from a single-node setup to a multi-node setup, you might not see a linear increase in performance. You'll need to test and fine-tune this multi-node configuration just as you did with the single-node setup to ensure optimal performance.
- **Broaden infrastructure:** Success in AI often depends on the entire system's efficiency. Review your current technology platforms, networks, and storage solutions to ensure they can handle increased data volume, improve data timeliness, and reduce latency. This approach helps prevent future slowdowns and maximises the potential benefits of your data.
- **Refine and optimise the PoC solution:** Over time, you'll gain insights that can enhance the AI solution you've deployed. You can improve aspects like data curation and labelling, and you may want to explore, train, and deploy new models that could deliver better results.
- **Apply the PoC to other business areas:** Consider how your PoC can be relevant to different parts of your organisation. For example, if you've implemented a predictive maintenance model in one area of manufacturing, think about how it could be expanded to other sections. A portfolio approach can help you manage this extension effectively and ensure a wider user base benefit.
- **Prepare for ongoing management and operations:** Many AI applications require real-time performance rather than offline processing. Additionally, your models will likely need to be retrained and updated periodically to stay efficient.

By following these steps, you can create a solid foundation for your AI initiatives and help them thrive as they grow.

Factors affecting AI PoC development costs

Several factors come into play when considering the costs associated with developing an AI Proof of Concept. Below, we summarise the key components affecting AI PoC development costs.

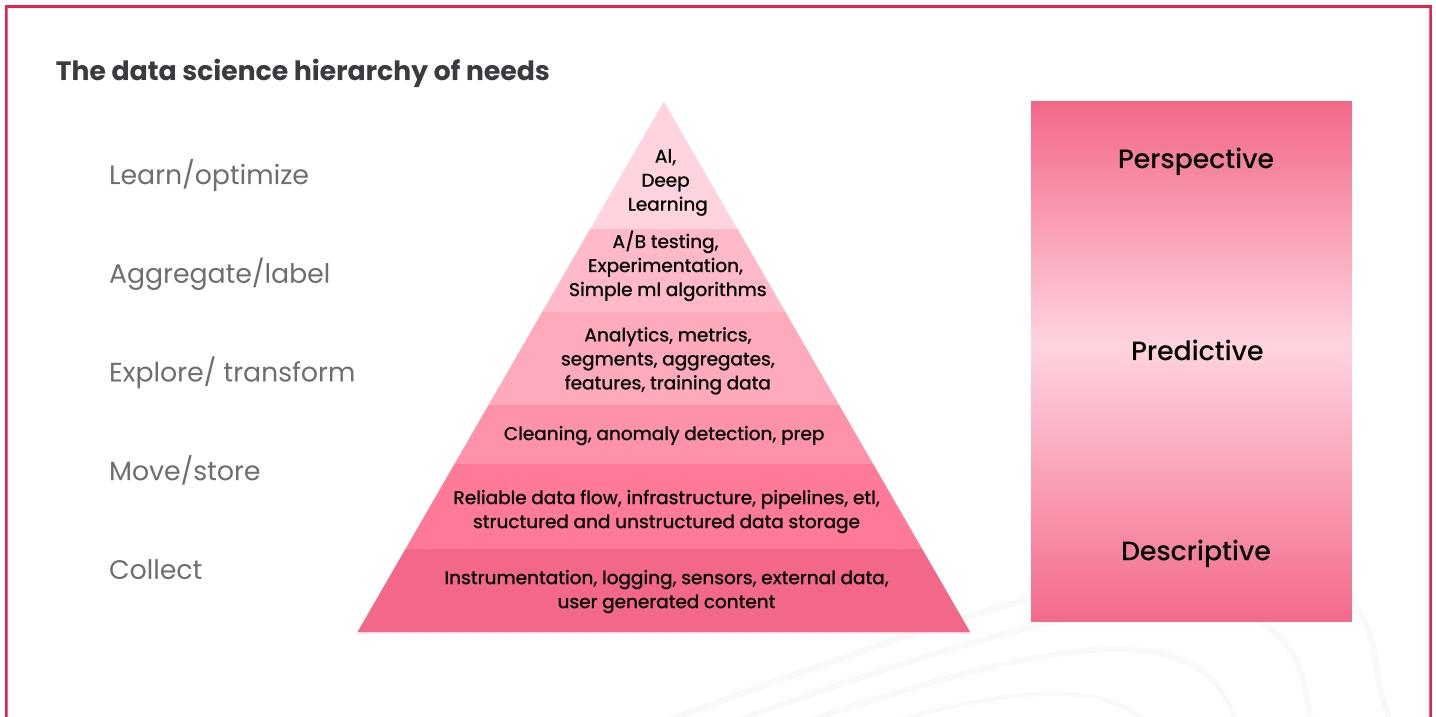
1. The project's scope

The scale and complexity of the problem at hand can greatly impact your overall budget. More complex issues often require sophisticated algorithms and more extensive data analysis, which can lead to increased costs. For example, a project that predicts customer behaviour might need a more advanced machine learning model than a simple classification task.

According to a study by McKinsey, companies can expect to spend anywhere from \$200,000 to \$1 million on AI initiatives, with complexity being a major cost driver.

2. Data availability and quality

The accessibility and quality of data are critical to the success of any AI project. If your organisation has ample, clean, and relevant data readily available, the costs associated with data preprocessing will be minimal. However, if data needs to be collected, cleaned, and formatted before it can be used, those additional tasks can inflate costs significantly.



*Gartner suggests that organisations may spend **up to 80% of their AI project budget** on data preparation and management, underscoring the importance of this factor.*

3. Technology stack and infrastructure

The choice of technology stack and infrastructure also plays an important role in determining costs. Cloud services can offer flexibility and scalability but may introduce ongoing expenses. For example, using platforms like AWS or Google Cloud leads to variable costs based on usage.

AWS currently provides eight cloud storage services, including Amazon S3 (object storage), Amazon EBS (EC2 block storage volumes), and Amazon EFS (file storage system). Each cloud storage service comes with its own pricing. For example, Amazon S3 pricing starts at \$0.023 per GB per month.

S3 Standard - General purpose storage for any type of data, typically used for frequently accessed data	
First 50 TB / Month	\$0.023 per GB
Next 450 TB / Month	\$0.022 per GB
Over 500 TB / Month	\$0.021 per GB
S3 Intelligent - Tiering * - Automatic cost savings for data with unknown or changing access patterns	
Monitoring and Automation, All Storage / Month (Objects > 128 KB)	\$0.0025 per 1,000 objects
Frequent Access Tier, First 50 TB / Month	\$0.023 per GB
Frequent Access Tier, Next 450 TB / Month	\$0.022 per GB
Frequent Access Tier, Over 500 TB / Month	\$0.021 per GB
Infrequent Access Tier, All Storage / Month	\$0.0125 per GB
Archive Instant Access Tier, All Storage / Month	\$0.004 per GB
S3 Intelligent - Tiering * - Optional asynchronous Archive Access tiers	
Archive Access Tier, All Storage / Month	\$0.0036 per GB
Deep Archive Access Tier, All Storage / Month	\$0.00099 per GB
S3 Standard - Infrequent Access ** - For long lived but infrequently accessed data that needs millisecond access	
All Storage / Month	\$0.0125 per GB
S3 One Zone - Infrequent Access ** - For re-createable infrequently accessed data that needs millisecond access	
All Storage / Month	\$0.01 per GB
S3 Glacier Instant Retrieval *** - For long-lived archive data accessed once a quarter with instant retrieval in milliseconds	
All Storage / Month	\$0.004 per GB
S3 Glacier Flexible Retrieval (Formerly S3 Glacier) *** - For long-term backups and archives with retrieval option from 1 minute to 12 hours	
All Storage / Month	\$0.0036 per GB
S3 Glacier Deep Archive *** - For long-term data archiving that is accessed once or twice in a year and can be restored within 12 hours	
All Storage / Month	\$0.00099 per GB

On the other corner, Google Cloud Storage also offers several types of storage. But, again, the starting standard rate is slightly more affordable per GB than AWS cloud storage (\$0.020 vs \$0.023 per GB). See below:

Location	Standard storage (per GB per Month)	Nearline storage (per GB per Month)	Coldline storage (per GB per Month)	Archive storage (per GB per Month)
Iowa (us-central1)	\$0.020	\$0.010	\$0.004	\$0.0012
South Carolina (us-east1)	\$0.020	\$0.010	\$0.004	\$0.0012
Northern Virginia (us-east4)	\$0.023	\$0.013	\$0.006	\$0.0025
Columbus (us-east5)	\$0.020	\$0.010	\$0.004	\$0.0012
Oregon (us-west1)	\$0.020	\$0.010	\$0.004	\$0.0012
Los Angeles (us-west2)	\$0.023	\$0.016	\$0.007	\$0.0025
Salt Lake City (us-west3)	\$0.023	\$0.016	\$0.007	\$0.0025
Las Vegas (us-west4)	\$0.023	\$0.013	\$0.006	\$0.0025
Dallas (us-south1)	\$0.020	\$0.010	\$0.004	\$0.0012
Montréal (northamerica-northeast1)	\$0.023	\$0.013	\$0.007	\$0.0025
Toronto (northamerica-northeast2)	\$0.023	\$0.013	\$0.007	\$0.0025



SERVICE	AWS PRICING	GOOGLE CLOUD PRICING
Compute (Virtual Machines)	EC2: Starting from \$0.0116 per hour (t3.micro)	Compute Engine: Starting from \$0.0100 per hour (f1-micro)
Storage	S3: Starting from \$0.023 per GB (Standard)	Cloud Storage: Starting from \$0.020 per GB (Standard)
Databases	RDS: Starting from \$0.018 per hour (db.t3.micro)	Cloud SQL: Starting from \$0.017 per hour (db-f1-micro)
Load Balancing	Application Load Balancer: \$0.008 per hour + \$0.008 per LCU	Cloud Load Balancing: \$0.008 per hour + \$0.0075 per GB served
Networking (Data Transfer)	\$0.09 per GB (first 10 TB)	\$0.12 per GB (first 1 TB)
AI/ML Services	SageMaker: Starting from \$0.058 per hour (ml.t2.medium)	AI Platform: Starting from \$0.10 per hour (n1-standard-1)

*Deloitte highlights that businesses spend an average of **\$20,000 per month** on cloud services for AI workloads. This cost will vary depending on the specific requirements and the scale of your PoC.*

4. Team expertise and staffing

Finally, the expertise of your team can heavily influence development costs. Hiring skilled data scientists, engineers, or external consultants typically represents a significant portion of the budget.

*The average salary for a data scientist in the U.S. is approximately **\$113,000 per year**, according to the U.S. Bureau of Labor Statistics.*

Additionally, organisations may need to account for training costs to upskill existing staff or the costs associated with hiring niche talent for specific project needs.

Developing an AI Proof of Concept involves various cost-driving factors, including the complexity of the problem, the availability and quality of data, the technology stack and infrastructure, and the team's expertise. By carefully considering these elements, organisations can better estimate their budgets and allocate resources wisely for successful AI project outcomes.

Hidden costs to consider in AI PoC development

When embarking on an AI Proof of Concept, you have also to be aware of the hidden costs that can creep in during the development process. Understanding these costs will help you better plan your budget and avoid unpleasant surprises later on.

1. Time overruns and project scope creep

One of the most common issues in AI projects is time overruns, often caused by project scope creep, which occurs when additional features or requirements are added after the initial planning stage, leading to increased workloads and extended timelines.

*According to Project Management Institute, **43%** of projects experience scope creep, which can inflate costs significantly.*

When this happens in an AI PoC, teams may need to reallocate resources, prolonging development time and increasing labour costs.

2. Integration with existing systems

Integrating your PoC with existing legacy systems can also introduce unexpected costs. Legacy systems may not easily accommodate new AI solutions, leading to additional expenses for adaptation and integration.

*Gartner found that **up to 75%** of AI projects encounter issues related to integration with current systems.*

This can require hiring specialised consultants or developers familiar with both the new and existing technologies, further increasing the budget.

3. Post-PoC support and maintenance

Once your PoC has been successfully developed, ongoing support and maintenance will also incur costs. AI systems often require continuous monitoring, adjustments, and updates to ensure optimal performance.

*Accenture highlights that businesses spend, on average, **20% to 30%** of their initial AI project budget on maintenance and support.*

So, it is wise to factor in these ongoing costs when planning your overall budget.

Breaking down the costs

Creating an AI PoC involves multiple stages, each with its own associated costs. Here's a breakdown of the key elements.

1. Data collection and preparation



Cost range:
\$5,000 - \$50,000+



Key factors:
The type and amount of data required, its availability, and the cleaning process.

Data is the foundation of any AI project. The initial step in creating an AI PoC is collecting, cleaning, and preparing data for analysis. Depending on your organisation's existing data infrastructure, this could be a relatively simple process or a major undertaking. Costs will increase if your data is spread across multiple platforms or needs significant cleaning and structuring.

For organisations that lack sufficient internal data, purchasing third-party datasets or using external APIs can add to the expense. Additionally, privacy and compliance measures (such as GDPR) can further complicate and increase the cost of data preparation.

*Forrester Research highlights that **up to 80% of the effort in AI and machine learning projects is spent on data preparation, cleaning, and organisation.** Costs for this stage can range widely depending on data availability, quality, and security requirements.*

***Data-driven businesses** already managing large datasets may see lower data preparation costs, while those starting from scratch can incur expenses of **up to \$50,000+**.*

***Third-party data providers** can charge anywhere from **\$0.02 to \$0.50 per token**, making external data acquisition a significant portion of the cost.*

2. Model development



Cost range:
\$10,000 - \$100,000+



Key factors:
The complexity of the AI model and the development team's expertise.

The cost of developing an AI model can range dramatically based on its complexity. For example, a simple predictive model will cost far less than a custom-built deep learning algorithm.

The expertise of the AI engineers and data scientists involved also plays a major role in determining costs. Highly skilled professionals imply higher rates, but they can often build more efficient models.

For companies looking to develop AI PoCs internally, there are also costs related to the tools and platforms needed, such as cloud services (AWS, Google Cloud, Azure) or specialised software (TensorFlow, PyTorch).

Cognilytica breaks down costs based on model type:

- Depending on the use case and data volume, a **simple machine learning model** (such as logistic regression) can cost between **\$10,000 and \$30,000**.
- **Advanced AI models** involving deep learning or neural networks require more computational power and expertise, raising costs to upwards of **\$100,000**.

3. Infrastructure and tools



Cost range:
\$1,000 - \$30,000+



Key factors:
Cloud services, hardware, and software tools.

Many AI projects require significant computational power, especially those dealing with large datasets or complex models. Cloud services like Amazon Web Services (AWS), Microsoft Azure, and Google Cloud are commonly used to rent servers and storage on demand.

Infrastructure costs can quickly add up for organisations running AI experiments at scale. Compute-intensive models (such as neural networks) may require powerful GPUs or TPUs, which significantly increase cloud service costs or investment in on-premise hardware.

*According to McKinsey, **infrastructure-related expenses can consume up to 50% of an AI project budget** if not optimised. Using spot instances and auto-scaling can reduce this cost, but for real-time AI applications, maintaining server uptime can raise expenses significantly.*

Google Cloud's AI Infrastructure offers discounts for reserved and committed use, allowing businesses to reduce long-term cloud costs, making infrastructure more affordable depending on the AI project size.

4. Testing and validation



Cost range:
\$5,000 - \$25,000+



Key factors:
Model evaluation, optimisation, and iterative testing.

Once the model is developed, it must be tested to ensure it meets the defined objectives, which often include the accuracy of predictions, ensuring the model is not biased, and refining the performance based on feedback. Testing and validation processes can become too resourceful, especially if you need to adjust and refine the model multiple times. Optimisation of the model might require repeated training cycles, adding to both time and cloud costs.

Testing and refining an AI model is often iterative and can become costly as the number of experiments increases.

*IDC estimates that large organisations spend between **\$5,000 and \$25,000** on testing and validating models depending on the complexity and the number of iterations required.*

*A Stanford University report on AI accuracy found that repeated testing for refining AI models can add **10–20%** to the total project cost, especially in industries requiring high precision (such as healthcare or finance).*

5. Team costs and collaboration



Cost range:
Highly variable



Key factors:
Size and expertise of the team, time commitment.

A well-rounded team is inevitable for creating an AI PoC. You will need data scientists, AI/ML engineers, domain experts, and possibly a project manager to keep things on track. Team costs vary based on location, expertise, and the time needed to complete the project.

For example, the average annual salary of an experienced AI engineer in the U.S. is over \$120,000, but this number can vary depending on geography and demand.

In some cases, organisations may also choose to work with external consultants or agencies to develop their AI PoC. In this case, you may expect higher costs; however, you may reduce project risk due to external niche expertise.

According to O'Reilly's AI Adoption in the Enterprise report, staffing is among the highest cost drivers for AI development projects. Salaries for data scientists, machine learning engineers, and project managers vary widely, but here are some estimates:

U.S. data scientists command an average salary of around **\$122,000 per year**, while machine learning engineers earn slightly more, averaging around **\$136,000 annually**.

Hiring freelance AI specialists can cost around **\$100 – \$250 per hour**, depending on the expertise required.

A Deloitte report highlights that specialised AI expertise costs even more in high-demand markets where salaries can reach **\$200,000 per year** for AI roles.

How to minimise AI PoC development costs

One effective way to manage costs is to start with a narrow, well-defined use case. Focusing on a small, testable problem allows you to limit the scope of your project and reduce initial expenses. Once the PoC is successful, you can gradually scale it up, addressing more complex challenges as resources and understanding grow. This approach helps to minimise risk and manage the overall budget more efficiently.

1. Use existing AI frameworks and open-source tools

Using existing AI frameworks and open-source tools can greatly reduce development costs. Many pre-built libraries and frameworks, such as TensorFlow or PyTorch, are available for free and can save substantial time and effort in development.

Research from Forrester shows that companies using **open-source solutions can reduce their software development costs by as much as 30%**.

2. Outsourcing vs. In-house development

Finally, consider the costs and benefits of outsourcing versus in-house development. While hiring an external team may initially seem more expensive, it can often be more cost-effective in the long run, especially if your internal team lacks the necessary expertise.

Deloitte found that **59%** of companies choose to outsource at least one function of their AI projects to reduce costs.

Weighing the pros and cons of each approach based on your organisation's specific needs and capabilities can help you make an informed decision. By being proactive in your planning and resource allocation, you can increase the likelihood of a successful AI project while keeping costs in check.

Real-world examples of AI PoC projects

When it comes to budgeting for an AI proof of concept, there's quite a range of costs depending on various factors, including the project's scope and specific requirements. Here are some real-world examples to provide insight into what you might expect when planning your own AI PoC:

AI-powered PoC solution for children's academic success empowerment

Tech stack:

Python, Dialogflow, Llama 3, TensorFlow, Flask, AWS

Altamira delivered an AI-powered education solution that helps students regulate emotions and provides learning strategies for their well-being while learning. The app should take the place of the executive functioning coach to help students navigate their education.

Traditionally, obtaining a neuropsychological evaluation for a child requires an in-person visit, often entailing long waits for appointments and the results. This process imposed unnecessary stress on children and delayed the implementation of required interventions.

The need for a more efficient method was evident, so the Client sought a tech partner to handle the custom build, which required a lot of nuances due to the highly regulated education industry. The Client's requirements included the following:

- < Develop a Proof of Concept to demonstrate the viability of a digital neuropsychological assessment tool.
- < Implement a subscription model to ensure a sustainable revenue stream.
- < Promote awareness by educating students and families about the tool, which aims to improve the experience of managing educational challenges.
- < Increase accessibility by making executive tutoring and coaching more available than traditional methods in a therapist's office.

The implementation process involved several phases needed to develop a high-performing, user-friendly, and compliant platform, covering:

- < Discovery stage
- < Web development of the MVP
- < Mobile development of the MVP
- < AI features integration
- < Maintenance and continuous improvement

Core features:

You should also consider explainability, which relates to the quality of decision-making. The criteria for explainability include:

- **Neuropsych testing:** Using AI and advanced analytics, our platform conducts assessments quickly and impartially to identify cognitive dysfunctions while enabling continuous monitoring of changes in cognitive abilities.
- **Coaching:** Based on the test outcomes, personalised recommendations are generated to improve executive functions necessary for learning, working, and daily life management.
- **Learning strategies:** Our platform offers learning techniques to assist students in tackling various tasks, thereby optimising their educational outcomes.
- **Integrations with LLMs:** Initially compatible with Google Classrooms, our roadmap includes integration with other major LMS platforms such as Schoology, Blackbaud, Black Board, Infinite Campus, and Canvas.
- **Parent dashboard:** The interface provides parents with insights into their child's progress and offers access to coaching suggestions.

As a result, the solution was successfully deployed, profoundly transforming the educational and developmental trajectory for kids, parents, and educators.




AI car damage recognition

Tech stack:

React, Python, OpenCV, Pillow, TensorFlow, EfficientNet, ResNet, Mask R-CNN, ChatGPT, Django, Azure

Altamira's AI-powered car damage detection solution uses photos captured with a smartphone to instantly detect dents, scratches, and minor body defects in vehicles. This revolutionary technology makes it possible to activate insurance policies on the spot without any need for expert assessment.



You can overcome the following challenges with our AI car damage detection solution:

- < **Personalised customer experience:** Generate personally adjusted prices based on the car owner's driving patterns, vehicle's condition, speed patterns, traffic environment, and the number of miles driven.
- < **Better risk assessments:** Make more informed and accurate decisions about the risks of insuring different vehicles.
- < **Faster incident investigation:** Assess the damages remotely, allowing quicker processing of insurance claims.
- < **Exact vehicle damage detection using AI:** Exactly evaluate the vehicle condition based on gathered and processed data generated by internal car sensors and get accurate calculations.
- < **Detection of fraudulent cases:** By providing objective, consistent, and accurate information about the condition of a vehicle, this technology can help insurers identify potentially fraudulent cases.

With our AI PoC solution you can:



Recognise a car

The system automatically assesses the car in the photo, its colour, model, and angle while detecting any signs of car damage. It also recognises whether the image was modified to prevent fraud.



Know the severity of the damage

Damage recognition contributes to rapid assessment of the car damage level and the location of the damaged components. It also identifies whether the car is damaged due to a car incident or merely surface-level dirt.



Get a repair estimation cost

By analysing data from a large number of repairs, a car damage detection solution powered by machine learning identifies patterns and trends in the cost of different types of repairs. So that the repair costs are estimated accurately, and insurance policies are priced reasonably.

Pro tips for advanced practitioners

For those who are looking to streamline their AI POC development, consider the following strategies:

1. Optimise data preprocessing with automation

If your team is experienced in handling large datasets, implementing automated data preprocessing pipelines can save significant time and reduce manual intervention. Tools such as Apache NiFi or Azure Data Factory allow you to automate data collection, cleaning, and transformation from various sources.

*According to a Forrester Consulting study, automating data pipelines with tools like **Apache NiFi** or **Azure Data Factory** can reduce manual labour and cut costs by up to 20%.*

2. Use pre-trained models

Instead of building every model from scratch, you can save time and costs by using pre-trained models and fine-tuning them to your specific needs. OpenAI's GPT models, Google's BERT, and many others offer robust pre-trained frameworks that can be adapted for particular use cases.

*Using **pre-trained models** from sources such as OpenAI, Hugging Face, or Google's TensorFlow Hub can reduce development time and costs by up to **40%**, especially when fine-tuning a pre-existing model.*

3. Experiment with open-source tools

For those looking to avoid licensing costs, numerous open-source tools can perform at an enterprise level. Platforms like TensorFlow, PyTorch, and Scikit-learn offer powerful capabilities without the associated expenses of proprietary software.

*Using open-source tools can cut licensing fees by **\$10,000 to \$50,000 annually**, depending on the software alternatives chosen.*

4. Cost-effective cloud usage

When using cloud platforms for AI development, consider using reserved instances, spot instances, or auto-scaling solutions to optimise your infrastructure costs. Many cloud providers offer discounts for long-term usage commitments or allow you to take advantage of unused server capacity at reduced rates.

*Gartner reports that using **spot instances** and **auto-scaling** can save organisations up to **60% on cloud infrastructure** costs during AI model training.*

5. Continuous model monitoring

Post-deployment, it is important to monitor AI models in real-time to ensure they continue to perform as expected. Consider implementing MLOps practices that automate the monitoring and retraining of models, which can prevent performance degradation over time and reduce maintenance costs.

Conclusion

The costs of creating an AI PoC can vary based on data, infrastructure, team, and model complexity. Organisations should allocate budgets carefully, considering upfront costs and potential long-term savings from automation, cloud optimisations, and pre-existing models. By using data-driven insights and optimising each stage of development, businesses can create AI PoCs that are cost-effective but efficient.

The traditional “one size fits all” model is no longer suitable. Today, the faster a company can implement a solution and demonstrate its value, the quicker it can generate revenue from adopted solutions.

How Altamira can help you

Altamira is a global digital transformation partner, helping organisations scale faster and more sustainably than anyone else.

With market-leading capabilities across all aspects of product and technology development, we make a difference where it matters, when it matters. We provide the following services:

- < Mobile and web product development and delivery
- < Fast and efficient scaling of your in-house development team
- < Boosted time to market
- < Rapid MVP creation and deployment
- < Expert consulting and technical auditing

Our compressed SDLC toolkits and use of ready-to-use components and prototypes enable us to deliver MVP/POC solutions faster, getting your product to market quickly. The reduced time to market means you can start generating revenue sooner, accelerating your return on investment and driving business growth.

Comprehensive database of ready-to-use components

Our extensive database of pre-built components allows us to rapidly assemble and customise AI solutions tailored to your needs. This reduces development time and costs while maintaining high quality.

Rich repository of prototypes

With a vast library of prototypes at our disposal, we can quickly create working models of your application. We help you to visualise the end product early in the development cycle and make informed decisions.

Expertise in non-functional and security requirements

Our team has deep expertise in addressing non-functional requirements such as performance, scalability, and maintainability. We ensure your software runs smoothly and efficiently, even under heavy loads.

Artificial intelligence consulting services

Our AI consulting services stand for strategy, governance, security, and implementation, driving your AI initiatives from concept to reality.

- < Quality of the data to be used within AI/ML solutions
- < Data annotations readiness
- < Feasibility of the pre-built AI/ML models
- < Envisioned business goals and their translation into technical solutions
- < Viability of the solutions, risks, and tradeoffs

AI data and analytics services

- < Scaling your ML initiatives, managing increased data volumes, model challenges, and computational resources as operations grow
- < Streamlining model deployment processes for faster time-to-market
- < Implementing MLOps practices for the reliability and governance of ML models
- < Automating repetitive tasks and workflows
- < Optimising costs associated with model development, deployment, and maintenance
- < Implementing data management practices for quality, integrity, and security
- < Continuous monitoring and retraining of ML models through MLOps

Large language model development and consulting

We create tailored LLMs using foundational models, such as ChatGPT, Gemini, Llama, Claude, and Mistral, and smoothly integrate them into your ecosystem.

- < User workflow design and project roadmap development
- < Dataset preparation
- < Integration of existing cloud-based AI solutions from AWS, Azure, and OpenAI
- < Proof of Concept development
- < Proprietary data overview
- < Fine-tuning of large language models
- < API integration
- < Modernisation of data architecture
- < Automation of large language models

Altamira's Chief Delivery Officer Nataliia Semak about the AI Proof of Concept process:

"At Altamira, we approach the AI Proof of Concept process with a focus on **reducing risks, confirming business value, and ensuring that proposed solutions are technically viable** before progressing to full-scale development. Given that AI projects can be heavy-duty, it's important to evaluate their feasibility and potential return on investment beforehand. A well-structured PoC helps us test key assumptions, identify technical challenges, and recognize potential risks, ensuring alignment with business goals.

Typically **completed in about 4 weeks, with costs ranging from \$15,000 to \$20,000** depending on the complexity of the project, our PoC process involves a tailored team that may include data annotators, engineers, data scientists, and AI/ML specialists.

A PoC becomes particularly vital when there are uncertainties about data availability, the expected business impact of the AI solution, or the risks related to a full-scale rollout. **By addressing these areas early, we can help prevent costly missteps in later stages of development.**

For clients looking to develop an AI PoC, **we recommend starting by clearly defining the business problem, establishing measurable success criteria, and focusing on a small, manageable project scope.** This strategy streamlines the process and provides valuable insights for future AI initiatives."

Have a question? Let's talk.

We love to help our clients be successful!

Email: sales@altamira.ai

www.altamira.ai