



From Data to Action

ARTIFICIAL INTELLIGENCE AND
MACHINE LEARNING SOLUTIONS

AI is here to transform organizations. So are we.

Artificial intelligence — or AI — has arrived, and its potential to transform organizations large and small can be seen everywhere we look, from market research to finance, insurance to logistics. But getting the most from AI requires a big effort to turn raw data into usable, accurate and actionable insights.

That's where Novacene can help. We automate the process needed to go from data to business value so that your organization and your customers can benefit from the full potential of AI.

www.novacene.ai

On this Booklet

What Does Novacene Do?

A description of our products and services

What is AI?

Straightforward descriptions of common AI terms

AI Use Cases

Examples to help spark ideas on how to use AI

How to Implement AI

How to discover, measure and implement AI-driven initiatives

Let's Work Together

A description of our process to help you get the most from AI



What Does Novacene Do?

We help organizations turn raw data to actionable insight. We do this by providing a self-serve platform and consulting services.

The Platform

The NovaceneAI Platform puts the power of AI and machine learning in the hands of non-technical business users. With the NovaceneAI Platform, you can:

- ✓ Upload data files or pull from existing data sources
- ✓ Make predictions using built-in models
- ✓ Train your own custom models
- ✓ Visualize data with interactive widgets
- ✓ Export predictions

The platform's functionality can be accessed through an intuitive web interface, or programmatically through an API.

Consulting Services

We provide a variety of data-focused services tailored to your specific needs. Our consulting services include:

- ✓ Machine learning model development, maintenance and support
- ✓ Setup, configuration and maintenance of cloud-based data science environments
- ✓ General data science consulting
- ✓ AI-driven value creation assessments, and use case identification





What is AI?

The technical and scientific nature of artificial intelligence technologies make AI an intimidating topic for most people. And the hype around what AI can do blurs the lines between reality and science fiction.

It doesn't have to be like this.

What is AI

While you might think that AI is too “bleeding edge” for your organization, it is very likely that you are already using AI in one way or another, at work and at home. Worry no more: let these descriptions guide you.

Artificial Intelligence (AI)

AI refers to a computer’s ability to perform tasks that resemble humans’ capabilities to perceive the world around them and to make decisions based on that perception. For example, AI enables computers to:

- Read documents and understand topics
- See objects in an image and categorize them
- Hear and understand instructions
- and more

Sensory capabilities like *seeing* and *hearing* are possible thanks to sensors like cameras and microphones. Meanwhile, Machine Learning facilitates the cognitive capabilities like *understanding* and *categorizing*.

Machine Learning (ML)

Machine Learning consists of a set of statistical methods that make it possible to predict an outcome. The larger the amount of data used to predict the outcome, the higher the likelihood that the prediction will be correct. Like humans, machines learn to make better decisions as their exposure to data increases.

Supervised Machine Learning

Supervised learning refers to the way in which a machine learning model is trained using labeled data.

For example, let’s say you are trying to automate the verification of user accounts to determine whether they are real or fake. The labeled data could be a spreadsheet with columns for the different attributes of an account and a column containing the corresponding label *real* or *fake*.

The machine learning algorithm would create a mathematical representation that reflects the relationships in the data, and use that representation to predict the labels for cases that it hasn’t encountered before, effectively predicting whether a new account will be real or fake.

Unsupervised Machine Learning

Unsupervised learning is another method used by computers to predict outcomes. This method does not require labeled data, however. Instead, the algorithm uses statistical methods to make predictions. For example, an unsupervised learning model could be used to categorize responses to the question “*Which issues drive your voting preferences?*” into categories such as *the economy*, *the environment*, *healthcare*, etc. In this case, the computer performs a few tasks: first, it converts the responses to a numerical representation. Then, it plots these numbers on a graph and groups nearby points together into clusters. These clusters represent the different categories.

Predictive Analytics

In the context of AI, there is a continuum of ways in which analytics can be visualized. You’ll hear terms like descriptive, diagnostic, predictive and prescriptive analytics. The first two refer to the ability to *describe* events that took place in the past, and the ability to *diagnose* patterns and relationships in the data related to those events. The last two provide *predictions* about future outcomes and *prescribe* actions to take to achieve future goals.



AI Use Cases

Here is a little inspiration to help you imagine how you could use AI to unleash the power of human + machine collaboration in your organization

AI business applications are here to stay. Organizations are using AI to improve operations, increase value to customers and reduce costs. But it can be hard to get started implementing AI when opportunities aren't clear. Here are a few real-world examples to help spark ideas you can try in your organization.

Prioritizing prospects

Digital marketing tools provide users with various reports about the performance of their campaigns. These reports include many key performance indicators and qualitative data, like who opened an email, who clicked on a link, etc. These tools also calculate and assign a *quality score* to each contact based on their behaviour. Customer Relationship Management systems do this too; these systems use the data housed in the system to score leads and surface the most qualified prospects. Sales and marketing teams could take this one step further. They could combine marketing, sales and inventory information and train a machine learning model that shows them who might buy what and when, enabling them to tune their marketing messaging to maximize conversions.

Automating the verification of information

Many organizations rely on humans to review and categorize information. Take for example a market research firm that automatically collects web content related to their area of research. A team of human verifiers then combs through the data and flags the useful information. This job becomes increasingly unsustainable as the team gets buried under a growing pile of records. A classification model using historical data could help prioritize the records that are more likely to be useful and potentially save the team thousands of hours per year.

Digitizing paper-based information

A lot of important information is trapped on paper and many scanned documents aren't searchable. Organizations can digitize this information using commercial or open source products.

Textract, one open source option, describes the problem on their website as follows: *"As undesirable as it might be, more often than not there is extremely useful information embedded in Word documents, PowerPoint presentations, PDFs, etc—so-called 'dark data'—that would be valuable for further textual analysis and visualization."* Imagine being able to extract customer information from paper forms, or to analyze financial reports trapped in PDF tables.

Categorizing sentiment and opinions

People who analyze surveys spend a lot of time categorizing answers to open-ended questions. Not only a time-consuming issue, humans are prone to bias when interpreting ambiguous information. This becomes even more problematic as multiple people—each with their own biases—divvy up the work of categorizing a single set of responses. Topic classification algorithms can cluster responses into logical categories. Sentiment analysis tools can determine whether a response is positive, negative or neutral. Other natural language understanding methods can flag threatening, abusive or inappropriate comments. These and other text-analysis tools can dramatically reduce the time spent on these tasks and increase the quality of the analysis.

Automating first line of support

Chatbots are everywhere. These virtual agents use Natural Language Understanding and machine learning models to match customers' questions to existing answers. While the technology isn't perfect—chatbots still feel a bit clunky—they will become very good sooner than later. Eventually, these bots will provide a good-enough experience and answer most questions or route the more nuanced ones to the appropriate human agents.

Improving the user experience of marketplace applications

To a large extent, the defensibility of a marketplace-style application depends on its ability to create and sustain network effects. Further, what makes a marketplace sticky and valuable is the frequency with which both sides of the market come together in a transaction. Following common business models, the higher the number of transactions, the more revenue the marketplace generates.

Machine learning can help entice the demand side to transact. A machine learning model could present specific supply goods—for example, featured homes for sale—to users who are likely to value the amenities of such homes. Also, the site's search engine could rank results based on the likelihood that users performing the search would find the options appealing.

Case Study

A market research firm used machine learning to automate the verification of database records. The machine learning prediction model enabled the company to reassign staff from repetitive tasks to higher value-creating work.

Challenge

A market research firm in the Educational Technology space needed to increase its capacity to verify, approve or reject, tens of thousands of records collected as part of their market insights services. The verification was carried out by a team of verifiers who struggled to keep up as new records became ready for verification. The firm needed to prioritize the records and process the ones most likely to be valid first.

Solution


NovaceneAI worked with the client to review the data and develop an approach to meet the business objectives. Following the planning stage, NovaceneAI developed a machine learning classifier that took signals from the records and predicted whether a new record was valid or invalid. Following fine-tuning, the selected model achieved the performance requirements set by the business.

“*The prediction model built by NovaceneAI is **saving us 200 hours of work per month.**”*
Said Justin Ménard, Founder and CEO of LISTedTECH—*“We can now use that time to work on a backlog of tasks.”*”

Outcome

The classification model was delivered through the NovaceneAI Platform and it is accessible by the client 24/7. Non-technical staff from the market research firm can now classify new records and retrain the model with new data at the click of a button. The model is also used programmatically by the company's technical team to automate predictions on a daily basis.

Some definitions of AI are helpful but lengthy. This list is an attempt at making it easy for non-technical business leaders to understand basic concepts of AI by providing to-the-point explanations and examples.



Cognitive Computing

An AI area that encompasses methods for computers to digitize the physical world. These methods include primarily computer vision and speech recognition. For example, cognitive computing can be used to:

- Digitize handwriting using computer vision
- Transcribe audio using speech recognition
- Analyze physical behaviours using video image processing



Natural Language Processing (NLP)

NLP is an AI area dedicated to the understanding of human language in different contexts. In practical terms, NLP is often used to:

- Determine sentiment in customer feedback
- Detect threats or abusive language
- Understand the intent of an inquiry



Predictive Analytics

Some of the most common uses of analytics is to try to understand which events trigger certain outcomes. While this aims to explain the past and the present, predictive analytics are about predicting the future. Organizations can use predictive analytics to:

- Analyze historical trends to forecast future activity
- Determine data relationships, patterns and irregularities
- Model scenarios to detect issues and improvements

Combining Areas

Some of these areas can be combined. For example, an application could transcribe handwriting from a printed customer satisfaction form and determine whether the feedback is positive or negative. In this case, computer vision is used to digitize the handwriting and natural language processing to conduct sentiment analysis.



How to Implement AI

Meet Maria, our fictional startup CEO. Maria has been thinking that AI could create new value for her product and help differentiate it from her competitors'. But Maria is not a technical person, so she needs help understanding how AI can help take her product to the next level.

You don't know what you don't know

Maria keeps on top of the latest technology trends: blockchain, cryptocurrency, quantum computing, artificial intelligence, virtual and augmented reality, etc. From these many trends, she believes that AI has the potential to transform her business right now. She is interested in exploring further, but she doesn't know where to start. Some of the questions she's asking herself keep her from taking the next step. Questions like:

- How exactly could AI make my product better?
- Is implementing AI really worth the effort?
- If it is worth it, how do I implement it?

Each of these questions has a specific answer. Read on to see how Maria developed her understanding and positioned the implementation to succeed.

What exactly can AI do for her product?

Integrating AI could make a product or service immensely more valuable to the point of transforming a company's entire value proposition; in other cases the effort just isn't worth it. So the first step is to determine what exactly AI can do for her product. The AI definitions and use cases like the ones found in this booklet should help Maria connect the dots and picture how AI applies to her situation.

How much value will AI bring to her product?

At this point, Maria has determined which AI areas apply to her product. She is now asking herself if in fact AI will solve the problem she identified, and if so, how much of the problem it will solve. Thankfully, there is a way to answer these two questions without embarking on a costly implementation: meet the **proof of concept**.

Next up: the proof of concept

The indisputable way to answer whether the value brought by AI justifies the effort to implement it, is to set up a proof of concept (POC) and measure the results. The POC will help Maria answer two key questions:

- Can AI actually accomplish what she expects?
- What are the measurable benefits that AI can bring to her product?

The first question is about feasibility and the answer should be a clear-cut *yes* or *no*. The answer to the second question should be quantifiable (e.g. *implementing AI would reduce a person's workload by 20%*)

So, when commissioning the POC, Maria needs to make sure it will answer these two questions.

How does she go about implementing AI?

By now, Maria knows which AI areas will bring value and how much value each will bring. The next logical step is to determine whether the value justifies the cost to implement the solution. This implementation effort is sometimes referred to as operationalizing AI, and in Maria's situation, it means integrating the POC with the rest of her product.

Knowing the quantifiable benefits that AI will bring and the cost to implement these benefits will help her measure ROI, perform cost-benefit analysis and gain confidence to approve the work.

To implement AI, she will need:


- A work plan outlining tasks, timelines and costs
- A team to carry out the work

Assembling the right team

Even though Maria runs a tech startup, her development team doesn't have the experience to implement AI. She should engage an external AI firm to help create a realistic work plan, assess level of effort and costs, and determine the skill gap that will need to be bridged.

The path forward

As a startup, Maria's business can't afford to embark on costly implementations without a strong sense of certainty. Implementing this approach will help her minimize risk and increase her chances of arriving at a successful outcome. Integrating AI could be the single most valuable feature in her product roadmap.



Let's Work Together

At Novacene, we follow a four-step process to deliver our services. Depending on the stage your organization is at, we go through some or all of the steps below.

1 Discovery

Some of our clients heard that AI could help them solve problems, but are not quite sure about what to do next. During the Discovery phase, we work with our clients to understand what they are trying to accomplish and whether they have the data needed to solve their problem.

2 Solution Design

With a clear understanding of the problem and data, we move to design the solution. Solutions can use one or more components that interact to achieve a goal. These components can include machine learning models, application programming interfaces and more.

3 Proof-of-concept (POC)

Unlike straightforward technologies, AI has an element of uncertainty when it comes to assert whether an approach will solve a given problem. Even if it does solve it, the performance of the solution may not justify the effort. A POC provides a low-risk approach to measuring the effect of the solution and determine next steps.

4 Operationalization

POCs are typically functional but rudimentary implementations of a larger solution. POCs tend to live in isolation from live systems and are normally not suitable for production. Operationalizing means scaling the POC for production-grade performance and integrating it with other systems and processes.



Ready to take the next step?

We hope that this booklet helped increase your understanding of AI and to imagine how you can use it to improve processes and create more value for your customers.

You might be wondering what to do next. At Novacene, our mission is to enable business leaders to leverage the power of AI. We are here to support you and your organization as it begins to plan and incorporate AI into day-to-day operations.

Reach out to us for a consultation. We are here to help you get the most from AI.

Email us at info@novacene.ai

