



MICROSOFT AI PLATFORM

Build Intelligent Software

Artificial Intelligence
productivity for every
developer and every
scenario



With the Azure platform and productivity services, you can create the next generation of applications that span an intelligent cloud and an intelligent edge powered by AI.

Use a comprehensive set of flexible AI Services for any scenario, enterprise-grade AI Infrastructure that run AI workloads anywhere at scale, and modern AI Tools for developers and data scientists to create AI solutions easily and with the maximum productivity.

This paper provides a technical overview of Microsoft AI platform to help developers get a jumpstart to build innovative applications that augment human abilities and experiences.

Contents

| | |
|---------------------------------------|---|
| Introduction..... | 3 |
| Microsoft AI platform - Overview..... | 3 |
| Benefits of AI platform..... | 3 |
| AI platform stack..... | 4 |
| AI Services..... | 4 |
| AI Infrastructure..... | 6 |
| AI Tools..... | 7 |
| Conclusion..... | 9 |
| References..... | 9 |

© 2017 Microsoft Corporation. This document is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS SUMMARY. The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

Introduction

Vast amounts of data, faster processing power, and increasingly smarter algorithms are powering artificial intelligence (AI) applications and associated use cases across consumer, finance, healthcare, manufacturing, transportation & logistics, and government sectors around the world - enabling smarter & intelligent applications to speak, listen, and make decisions in unprecedented ways. As AI technologies and deployments sweep through virtually every industry, a wide range of use cases are beginning to illustrate the potential business opportunities, and inspire changes to existing business processes leading to newer business models.

Microsoft AI platform - Overview

The Microsoft AI platform offers a comprehensive set of flexible **AI Services**, enterprise-grade **AI Infrastructure** and modern **AI Tools** for developers and data scientists to create applications of the future.

AI platform consists of 3 core areas:

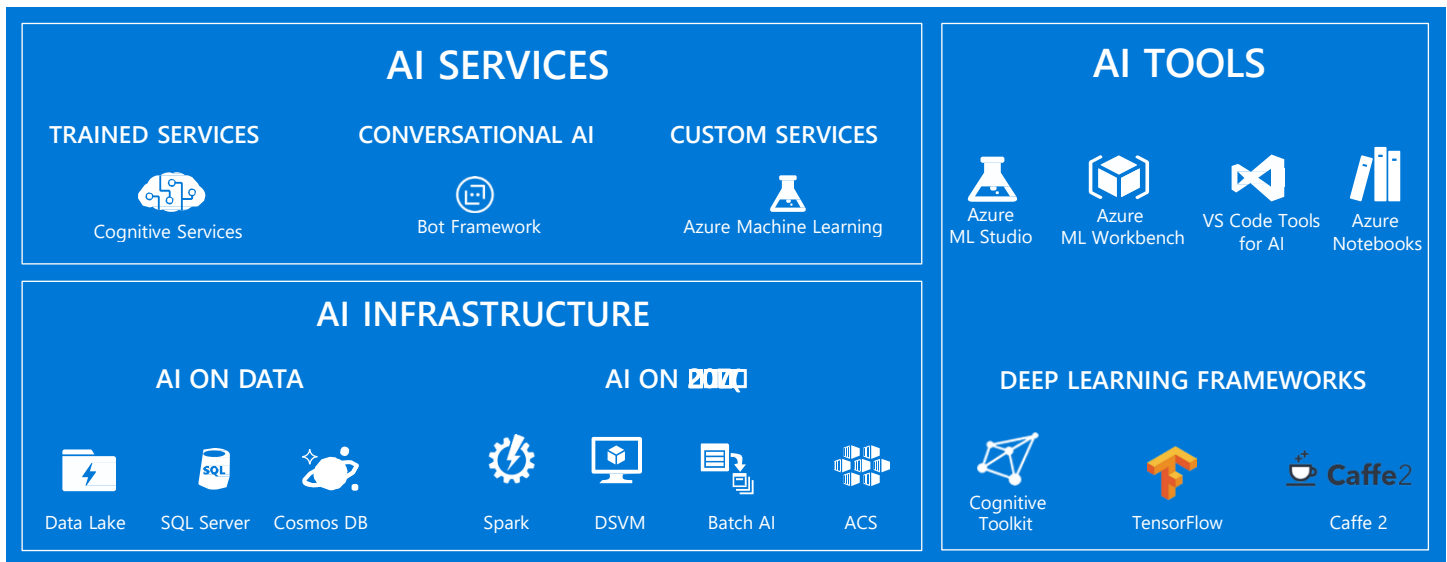
- **AI Services:** Developers can rapidly consume high-level “finished” services that accelerate development of AI solutions. Compose intelligent applications, customized to your organization’s availability, security, and compliance requirements.
- **AI Infrastructure:** Services and tools backed by a best-of-breed infrastructure with enterprise grade security, availability, compliance, and manageability. Harness the power of infinite scale infrastructure and integrated AI services.
- **AI Tools:** Leverage a set of comprehensive tools and frameworks to build, deploy, and operationalize AI products and services at scale. Use the extensive set of supported tools and IDEs of your choice and harness the intelligence with massive datasets through deep learning frameworks of your choice.

Benefits of AI platform

- The Microsoft AI platform offers finished AI services for rapid development, and provides data science tools to innovate and operationalize AI products and services at scale
- Easily customize your own models for unique use cases with easy-to-use customizable web services
- Rapidly compose intelligent applications with extensive APIs, customized to your enterprise’s security, compliance, availability, and SLA requirements
- Build immersive applications easily with intelligent features – such as emotion and sentiment detection, vision and speech recognition, language understanding, knowledge, and search – into your app, across devices such as iOS, Android, and Windows
- Leverage extensive deep learning frameworks of your choice - including Cognitive Toolkit, Caffe2, TensorFlow, Chainer, MxNet, Torch, Scikit-learn, and more
- Explore the extensive choice of IDE and data science tools – Azure ML Studio, Visual Studio, Azure ML Workbench, Jupyter Notebooks, PyCharm, or Juno
- Deploy your solutions on infrastructure that can virtually scale infinitely – all with enterprise grade security, compliance, availability, manageability including dev-ops capabilities such as Continuous Integration/Continuous Delivery (CI/CD) support for AI
- Create new immersive and integrated experiences - reach your users at scale. Easily build and deploy across channels including Facebook Messenger, Cortana, Slack, Skype, and Bing.

“AI is going to disrupt every single business app – whether an industry vertical like banking, retail and health care, or a horizontal business process like sales, marketing and customer support.”

- Harry Shum, Microsoft Executive VP, AI and Research



AI platform stack

Microsoft AI platform stack offers a rich set of interoperable services, APIs, libraries, frameworks and tools that developers can leverage to build smart applications.

AI Services

Compose intelligent applications, customized to your organization’s availability, security, and compliance requirements with a comprehensive set of flexible cloud AI Services.

Accelerate the development of AI solutions with high-level services. Use your preferred approach adapted to the scenario you are targeting with maximum productivity and reliability.

- **Cognitive Services:** Use AI to solve business problems. Infuse your apps, websites, and bots with intelligent algorithms to see, hear, speak, and understand natural methods of communication.
- **Bot Framework:** Accelerate development for conversational AI. Integrate seamlessly with Cortana, Office 365, Slack, Facebook Messenger, and more.
- **Azure Machine Learning:** Model AI algorithms and experiment with ease, and customize based on your requirements

Cognitive Services

Microsoft Cognitive Services expands on Microsoft’s evolving portfolio of machine learning APIs and enables developers to easily add intelligent features into their applications.

Cognitive Services are a set of APIs, SDKs, and services available to developers to make their applications more intelligent,

engaging, and discoverable and they let you build apps with powerful algorithms to see, hear, speak, understand, and interpret our needs using natural methods of communication, with just a few lines of code. Leverage customizable web services such as Custom Vision Service that can be trained to recognize specific content in imagery. Easily add intelligent features – such as emotion and sentiment detection, vision and speech recognition, language understanding, knowledge, and search – into your app, across devices such as iOS, Android, and Windows, keep improving, and are easy to set up.

Cognitive Services consist of the following services:

- **Vision:** State-of-the-art image processing algorithms help you moderate content automatically and build more personalized apps by returning smart insights
- **Speech:** Process spoken language in your applications
- **Language:** Allow your apps to process natural language, evaluate sentiment and topics, and learn how to recognize what users want
- **Knowledge:** Map complex information and data in-order to solve tasks such as intelligent recommendations and semantic search
- **Search:** Make your apps, webpages, and other experiences smarter and more engaging with the Bing Search APIs

Bot Framework

Think of a bot as an app that users interact with in a conversational way. Bots can communicate conversationally with text, cards, or speech. The Bot Framework enables you to build bots that support different types of interactions with users.

Bot Framework: Think of a bot as an app that users interact with in a conversational way. Bots can communicate conversationally with text, cards, or speech. The Bot Framework enables you to build bots that support different types of interactions with users.

You can design conversations in your bot to be freeform. Your bot can also have more guided interactions where it provides the user choices or actions. The conversation can use simple text strings or increasingly complex, rich cards that contain text, images, and action buttons. You can add natural language interactions, which let your users interact with your bots in a natural and expressive way.

A bot may be as simple as basic pattern matching with a response, or it may be a sophisticated weaving of artificial intelligence techniques with complex conversational state tracking and integration to existing business services.

The Microsoft Bot Framework makes it easy for you to create new experiences and reach your users at scale. Easily build and deploy across channels including Facebook Messenger, Cortana, Slack, Skype, and Bing.

You can build your bot with the Bot Builder SDK using C# or Node.js, or use the Azure Bot Service (currently in preview).


Add artificial intelligence to your bot with Cognitive Services. When you are ready to share your bot with the world, deploy it to a cloud service such as Microsoft Azure.


The Bot Framework is a platform for building, connecting, testing, and deploying powerful and intelligent bots. With support for .NET, Node.js, and REST, you can get the Bot Builder SDK and quickly start building bots with the Bot Framework. In addition, you can take advantage of Microsoft Cognitive Services to add smart features like natural language understanding, image recognition, speech, and more.

The **Azure Bot Service** provides an integrated environment purpose-built for bot development. You can write a bot, connect, test, deploy, and manage it from your web browser with no separate editor or source control required. For simple bots, you may not need to write code at all. It is powered by Microsoft Bot Framework and Azure Functions, which means that your bot will run in a server-less environment on Azure that will scale based upon demand.

Azure Machine Learning

Azure Machine Learning is a cloud predictive analytics service that makes it possible to quickly create and deploy predictive models as analytics solutions. The Machine Learning service is cloud-based, provides compute resource and memory flexibility, and eliminates setup and installation concerns because you can work through your web browser on any Internet-connected PC.





Azure Machine Learning service helps build, deploy and manage applications at scale. It helps boost productivity with agile development and enables you to begin building now with the tools and platforms you know.

Machine learning is considered a subcategory of artificial intelligence (AI). Forecasts or predictions from machine learning can make apps and devices smarter. For instance, you could build recommendation services - when you shop online, machine learning helps recommend other products you might like based on what you've purchased.

You can work from a ready-to-use library of algorithms, use them to create models on an internet-connected PC, and deploy your predictive solution quickly. Start from ready-to-use examples and solutions in the Cortana Intelligence Gallery.

Leverage the set of finished AI services to build immersive applications that use state of the art image processing with Deep Neural Networks (DNN) and explore the power of Natural Language Processing (NLP) capabilities for speech recognition. Use the extensive set of AI Tools supported to build rich immersive experiences.

AI Infrastructure

Leverage the power of virtually infinite scale AI infrastructure and integrated AI services.

AI Compute

Flexible compute services from virtually infinite scale to the edge

- **Spark on HDInsight:** Leverage Apache Spark in the cloud for mission critical deployments
- **Data Science VM:** Use friction-free data science environment that contains popular tools for data exploration, modeling and development activities
- **Batch AI Training:** Experience unlimited elastic scale-out deep learning. Perform massively parallel scale-out GPU enabled AI development.

- **Azure Container Service:** Deploy AI models with flexibility of containers and scale them out automatically with Kubernetes. Turn your AI models into web services using Docker containers. Auto scale and manage with Kubernetes.

Data Science VM (DSVM)

The Microsoft Data Science Virtual Machine (DSVM) is a powerful data science development environment that enables you to perform various data exploration and modeling tasks. The environment comes already built and bundled with several popular data analytics tools that make it easy to get started quickly with your analysis for On-premises, Cloud, or hybrid deployments.

You can use languages like R and Python to do your data analytics right on the DSVM. You can also leverage Jupyter Notebook that provides a powerful browser-based "IDE" for data exploration and modeling. You can use Python 2, Python 3 or R (both Open Source and the Microsoft R Server) in a Jupyter Notebook.

The DSVM works closely with many Azure services and can read and process data that is already stored on Azure, in Azure SQL Data Warehouse, Azure Data Lake, Azure Storage, or in Azure Cosmos DB. It can also leverage other analytics tools such as Azure Machine Learning and Azure Data Factory.

AI on data

AI enable your data platform

- **Data Lake:** Run data transformations and AI on petabyte-scale
- **SQL Server 2017:** Use R, python, and native machine learning in an industry leading SQL DB
- **Cosmos DB:** Integrate AI with a globally distributed multi-model DB storage

AI Tools

AI platform consists of comprehensive and productive tooling for AI coding and management. It enables developers to harness intelligence with massive datasets through tools and deep learning frameworks of your choice.

Coding and Management tools

AI platform provides a rich set of tools to simplify development:

- **Azure Machine Learning Studio:** Serverless collaborative drag-and-drop tool for graphical machine learning development
- **Azure Machine Learning Workbench:** Visual AI powered data wrangling, experimentation, and lifecycle management
- **Visual Studio Code Tools for AI:** Build, debug, test, and deploy AI with Visual Studio Code on Windows and Mac
- **Azure Notebooks:** Organize your datasets and Jupyter Notebooks in a centralized library for Data Science and Analysis


Aside from this, the platform supports several popular Open Source tools such as Jupyter Notebooks, PyCharm, and more.

Azure ML Studio

Azure Machine Learning Studio gives you an interactive, visual workspace to easily build, test, and iterate on a predictive analysis model. You drag-and-drop datasets and analysis modules onto an interactive canvas, connecting them together to form an experiment, which you run in Machine Learning Studio. To iterate on your model design, you edit the experiment, save a copy if desired, and run it again. When you're ready, you can convert your training experiment to a predictive experiment, and then publish it as a web service so that your model can be accessed by others.

Azure ML does more than just deploy a model - It automatically sets up the model to work with Azure's load balancing technology. This lets the model grow to handle cloud burst scenarios, scaling up to meet with use demands and shrinking when demand falls.

Azure ML studio also offers several standard templates - A machine learning template demonstrates the standard industry practices and common building blocks in building a machine learning solution for a specific domain, starting from data preparation, data processing, feature engineering, model training to model deployment.



The goal of the templates is to enable data scientists to quickly build and deploy custom machine learning solutions with Azure Machine Learning platform, and increase their productivity with a higher starting point. The template includes a collection of pre-configured Azure ML modules, as well as custom R scripts in the Execute R Script modules, to enable an end-to-end solution.

Azure ML Workbench

Workbench is visual AI powered data wrangling, experimentation, and lifecycle management tool. Tie it all together with Azure ML Workbench, that enables built-in data preparation that learns your data preparation steps as you perform them. Project management, run history, and notebook integration unleashes your productivity. Leverage the best open source frameworks such as TensorFlow, Cognitive Toolkit, Spark ML, Scikit-learn, and more.

VS Code Tools for AI

Build Deep Learning models easier, with Azure Machine Learning services built right in! Use Visual Studio Code Tools for AI to build, debug, test, and deploy AI on Windows and Mac for a seamless developer experience across desktop, cloud and edge. Develop deep learning models and call services straight from your favorite IDE.

Azure Notebooks

Leverage Azure Notebooks to organize your datasets and Jupyter Notebooks – all in one centralized location for your Data Science and Analysis. For instance, leverage Azure Notebooks to run negative matrix factorization (NMF) over large datasets easily and identify topics of interest on Twitter feeds.

Deep Learning Frameworks

AI platform stack supports an extensive array of deep learning frameworks – including Cognitive Toolkit, Caffe2, TensorFlow, Chainer, MxNet, Torch, Scikit-learn, and more. Deep learning is impacting everything from healthcare to transportation to manufacturing, and more. Companies are turning to deep learning to solve hard problems, like image classification, speech recognition, object recognition, and machine translation.

Deep neural networks (DNNs) are extraordinarily versatile artificial intelligence models that have achieved widespread use over the last five years. These neural networks excel at automated feature creation and processing of complex data types

like images, audio, and free-form text. Common business use cases for DNNs include:

- Determining whether an uploaded video, audio, or text file contains inappropriate content
- Inferring a user 's intent from their spoken or typed input
- Identifying objects or persons in a still image
- Translating speech or text between languages or modalities

Unfortunately, DNNs are also among the most time - and resource-intensive machine learning models. Whereas a trained linear regression model results can typically score input in negligible time, applying a DNN to a single file of interest may take hundreds or thousands of milliseconds -- a processing rate insufficient for some business needs.

To overcome the time complexity, DNNs can be applied in parallel – using a scalable fashion with Spark clusters. AI platform provides rich support for parallelism with Spark clusters. Leverage DNNs created with Cognitive Toolkit or TensorFlow, operationalized on Spark with Azure Data Lake as the store.

Cognitive Toolkit (CNTK)

Cognitive Toolkit will enable enterprise-ready, production-grade AI by allowing users to create, train, and evaluate their own neural networks that can then scale efficiently across multiple GPUs and multiple machines on massive data sets.

Cognitive Toolkit is a framework for describing learning machines. Although intended for neural networks, the learning machines are arbitrary in that the logic of the machine is described by a series of computational steps in a Computational Network.

CNTK can be included as a library in your Python, C#, or C++ programs. Additionally, you can use the CNTK model evaluation functionality from your Java program. With support for Keras, users will now benefit from the performance of CNTK without any changes to their existing Keras recipes.

Computational Network defines the function to be learned as a directed graph where each leaf node consists of an input value or parameter, and each non-leaf node represents a matrix or tensor operation upon its children. The beauty of Cognitive Toolkit is that once a computational network has been described, all the computation required to learn the network parameters are taken care of automatically. There is no need to derive gradients analytically or to code the interactions between variables for backpropagation.

Conclusion

Compose intelligent applications, customized to your organization's availability, security, and compliance requirements with Microsoft AI platform. With the Azure platform and productivity services, you can create the next generation of applications that span an intelligent cloud and an intelligent edge powered by AI.

Use a comprehensive set of flexible AI Services for any scenario, enterprise-grade AI Infrastructure that run AI workloads anywhere at scale, and modern AI Tools for developers and data scientists to create AI solutions easily and with the maximum productivity.

For more information and to learn more, refer to online training resources for AI Platform:

<https://azure.microsoft.com/en-us/training/learning-paths/azure-ai-developer>




Get Cloud AI Certified

Build expertise and advance your knowledge with Azure AI certification for Machine Learning.

<https://www.microsoft.com/en-us/learning/mcsa-machine-learning.aspx>

References

1. Microsoft Azure Notebooks: <https://notebooks.azure.com/>
2. Microsoft Cognitive Toolkit: <https://www.microsoft.com/en-us/cognitive-toolkit/>
3. Azure Machine Learning Studio: <https://azure.microsoft.com/en-us/services/machine-learning/>
4. Azure Machine Learning Workbench: <https://azure.microsoft.com/en-us/resources/videos/overview-of-ml/>
5. TensorFlow: <https://www.tensorflow.org/>
6. MxNet: <https://mxnet.incubator.apache.org/>
7. Caffe2: <https://caffe2.ai/>
8. PyCharm: <https://www.jetbrains.com/pycharm/>
9. Juno: <http://junolab.org/>
10. Keras: <https://keras.io/>



MICROSOFT AI PLATFORM

azure.microsoft.com/ai

