

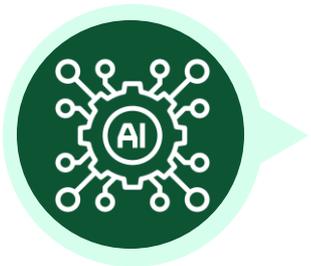
Getting Started with Data Extraction and AI





Poll:
Are you familiar with
Generative AI and use it?

What is Generative AI?



Most people are familiar with Generative AI (or GenAI) in the form of ChatGPT or other AI chatbots.

Put simply, GenAI are machine learning models that can **generate text, images, video, or other content** from end-user queries based on data they've been trained on.

Who are the **players** in the space?



For the sake of today's discussion we will be covering **Azure OpenAI** and **AWS Textract** Custom Queries in our demonstrations.



Both Cloud tools have **no user interface (UI)** to process documents in batch



So, how do we consume these Cloud services?

Intelligent Document Processing.



- IDP mimics the data processing employee to recognize, scan, and categorize documents.
- Relevant information is extracted using cognitive AI technologies like Large Language Models (LLMs) and neural transformers
- These pinpoint exact data points within the document
- IDP organizes the extracted data and presents it in a readily accessible format for routing, review, and approval within an ERP (Enterprise Resource Planning) system or applicable line of business system.

What's the Difference Between GenAI and Document AI?

Gen AI Use Cases



Azure OpenAI

- Use pre-existing LLM's like OpenAI, Claude, Llama, Grok
- Do NOT build models based on knowledgebase of samples and training data

Document AI Use Cases



- Must build models
- Must continue to train model based on variations in documents
- Building models takes time

What's the Difference Between GenAI and Document AI?

Gen AI Use Cases



Azure OpenAI

- Unstructured documents
- Extracting data from paragraphs of text
- Speed of deployment
- No uptraining
- Prompt engineering
 - Curate a prompt that works across all documents
- Less expensive (depends on model)

Document AI Use Cases

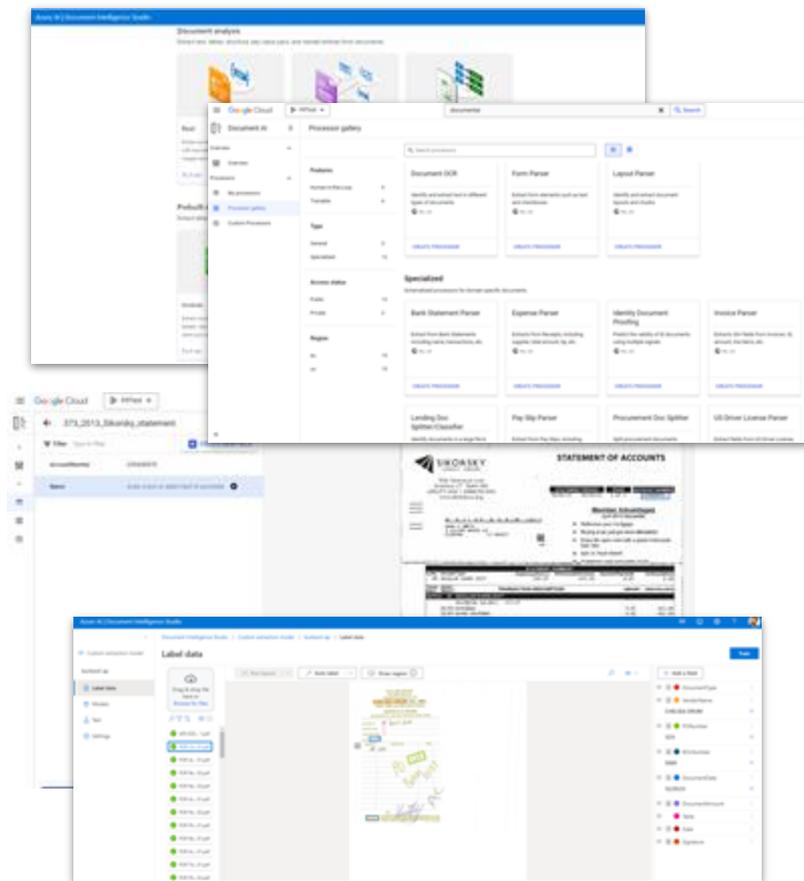


- Transactional documents
- Key-value pair extraction
- Prebuilt use case specific models (invoices, ID's)
- Supports uptraining
- Few shot historical predictions
- More expensive

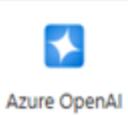
Prebuilt vs. Custom Models vs. Queries

The way you consume models impacts commercials and outcome

- **Prebuilt** - Small language models (SLM) for popular use cases
 - Returns a default set of extracted field names
 - Cannot uptrain
 - Less expensive than custom - \$0.01/page vs. \$0.03
 - No need to train knowledgebase and build models
- **Custom** - customer specific use case data models
 - Multiple document types in a batch can be recognized
 - Full uptraining and customization including tables
 - Does not support pagination of tables
 - Access to build models - share and import
 - Need to build new models for each change
- **Queries** (LLM's)
 - Azure OpenAi and AWS Textract Queries
 - Building adapters vs. autonomous
 - Great for unstructured documents where no key value pairs exist
 - General prompts that can extract data across diverse document set
 - Classify by asking the type of document



Poll:
What's your perception on
the cost to implement GenAI
in your business?
Very inexpensive or expensive?



Azure OpenAI - What Does it Cost?

What is Azure OpenAI compared to Azure AI Document Intelligence?

- Azure OpenAI is the generative ai offering for unstructured document extraction which differs from Azure AI Document Intelligence in that models are not pre-built, prompt engineering is used to structure common questions that extract the correct data.

How is pricing configured?

- **Pricing is based on input and output "tokens"**
- A token = ~4 characters
- Assumption below is 3,248 characters per page or 812 tokens/page

How does that compare with Azure AI Document Intelligence?

- \$0.01/page vs. \$0.00061/page - 93% difference



How much does it cost to extract 1 page?

- GPT-4o:
 - Input: (812\1,000,000 tokens x \$5 = \$0.00406)
 - Output: (812\1,000,000 tokens x \$15 = \$0.01218)
 - Total: \$0.01624 per page
- GPT-4o mini:
 - Input: (812\1,000,000 tokens x \$0.15 = \$0.00012)
 - Output: (812\1,000,000 tokens x \$0.60 = \$0.00049)
 - Total: \$0.00061 per page

AWS Textract - What Does it Cost?

AWS Textract pricing is much easier to understand

Consumption model by page

- \$0.015/page - up to 1M/month
- \$0.010/page - over 1M pages/month

AWS compared to Azure OpenAi

- \$0.015 vs. \$0.00061/page
- 96% difference

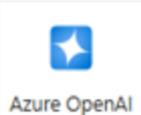
Analyze Document API		First million pages in a month	Over 1 million pages in a month
Queries	Per 1,000 Pages	\$15.00	\$10.00
Tables	Per 1,000 Pages	\$15.00	\$10.00
Tables + Queries	Per 1,000 Pages	\$20.00	\$15.00
Forms	Per 1,000 Pages	\$50.00	\$40.00
Forms + Queries	Per 1,000 Pages	\$55.00	\$45.00
Tables + Forms	Per 1,000 Pages	\$15.00 + \$50.00	\$10.00 + \$40.00
Tables + Forms + Queries	Per 1,000 Pages	\$70.00	\$55.00
Signatures	Per 1,000 Pages	\$3.50	\$1.40
Layout	Per 1,000 Pages	\$4.00	\$3.00

*Layout feature is included free of cost with any combination of Forms , Tables, and Queries

*Signatures feature is included free of cost with any combination of Forms , Tables, Queries, and Layout



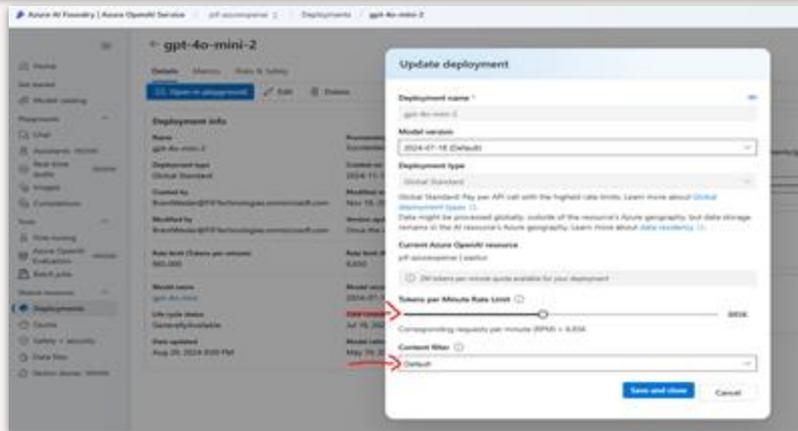
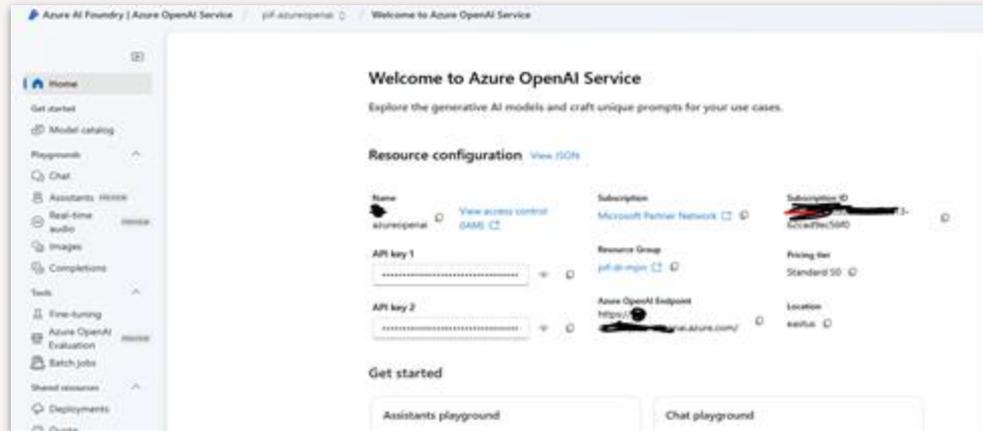
How to Configure Azure OpenAI and AWS



Azure OpenAI Configuration

Configuring Azure OpenAI Studio

- Using chat to test out queries
- Content filter settings
- Tokens per minute settings
- How do we configure Kodak to send and receive data from Azure OpenAI?
 - Query and key-value pair extraction at the same time
 - Modifying questions and adding new questions



AWS Textract Configuration

Custom Query Adapters vs. Textract Prompts

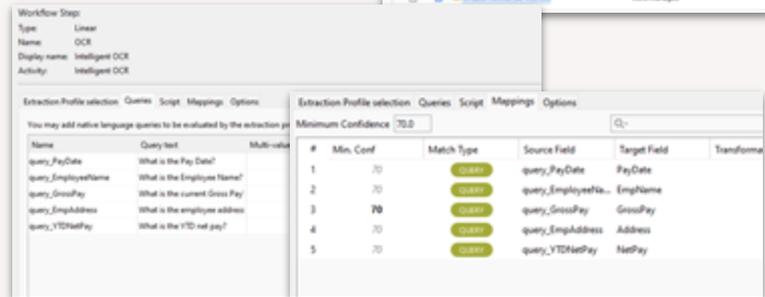
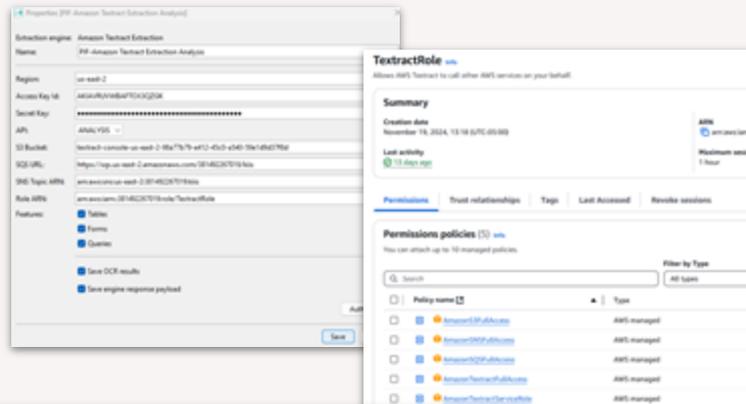
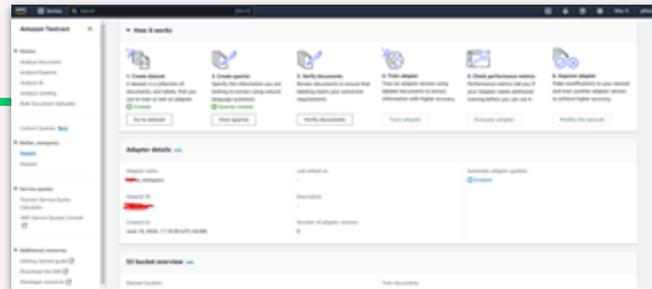
- Compiling adapters with defined document and data sets vs. autonomous queries

AWS Bedrock vs. AWS Textract Custom Queries

- **AWS Textract Custom Queries** is specialized for extracting specific data from documents
- **AWS Bedrock** provides a broader range of AI capabilities, making it suitable for more complex natural language processing tasks.

How do we configure Kodak to send and receive data from AWS?

- Query match type
- Query prompts



Discuss:

How does your organization stack rank the following when choosing to invest in new technology?

- Cost
- ROI
- Functionality
- Integration with existing systems

Document AI Classification & Extraction with Hypervisor Voting

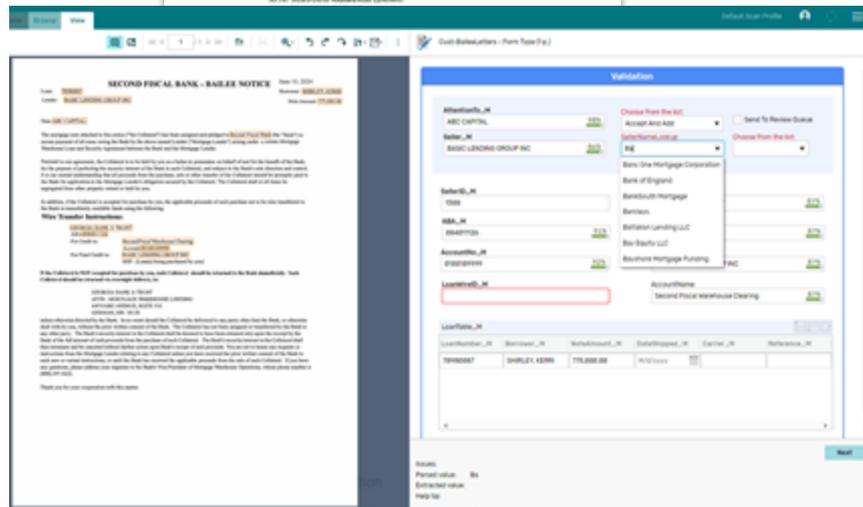
Use Case #1: Financial Services

The Challenge

- Getting both Azure and Google models to “line up”
- Normalizing extracted seller name to LMS seller name, and then populating seller ID and loan wire ID
- Creating a custom UI for triaging LMS lookups

The Solution

- Microsoft Azure AI Document Intelligence custom model accuracy was much better than Google
- Kodak Alaris Info Input batch capture “wrapper”
- Info Input Classify Designer - split on bailee first page
- Lots of Javascript logic to compare:
 - AlternativeSellerName - create alias'
 - Ability to Save and Add alias or use once
 - Based on ABA, account and seller ID populate the WireID
 - Verify the loan number extracted matches a valid loan number in LMS



Proprietary and confidential information. Do not distribute.



Microsoft Power Platform AI Builder & Encodian

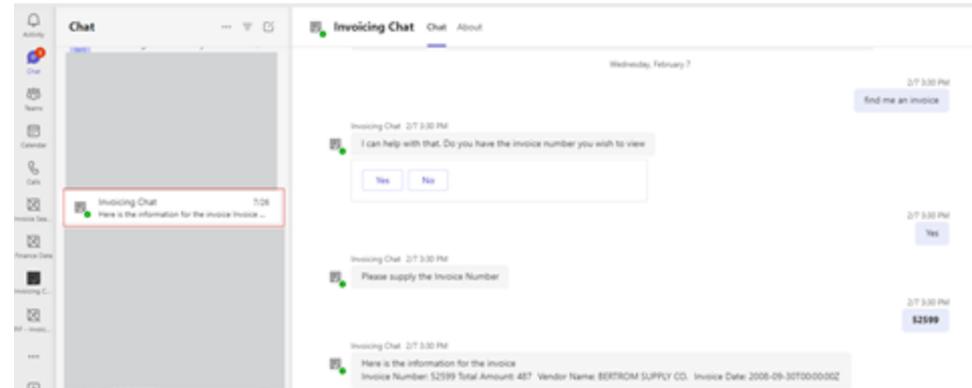
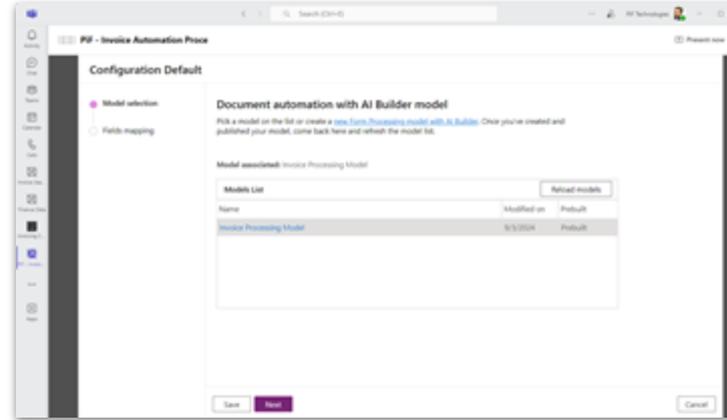
Use Case #2: Solar Energy

The Challenge

- Document library limits in Sharepoint
- Photo uploads from the field with offline access
- Integration into 3 ERP and CRM systems

The Solution

- Microsoft PowerPlatform - Using PowerApps, Power Automate and AI builder
- Power App Base Document Kit
- AI builder integration for invoice models
- Dataverse to manage all the data for Chat Ai prompting
- Encodian Filer, Trigr and Flowr to handle the document manipulation, storage, retention, searching and workflow



Doc AI Classification & RPA for Financial Statements

Use Case #4: Tax & Operating Statements

Profile of Customer

- Large hedge fund, publicly traded financial services company and their Boston based property management company

What was the pain?

- Receives tax documents and operating statements (balance sheets and income statements) from each property quarterly
- Documents are manually reviewed and manually entered into their financial system, line by line
- Large volume of tax and operating statements that come from 8,100 properties, all different layouts and formats.

Who were we selling to?

- CTO



Use Case #4: Tax & Operating Statements

The Challenge

- Each property operating statements are different formats
- Incoming paperwork may not have the correct documents needed to enter into financial system
- Tax documents while structured, come in as a packet and only certain pages of the tax packet need to be extracted

The Solution

- Qualify out pages through Azure classification model
- Azure AI Document Intelligence custom models to extract only those identified pages at \$0.03/page
- Azure AI custom extraction model annotating
- Alias learning based on the operating statement label
- UiPath Unattended and Attended bots for data entry
- UiPath Document Understanding cost of \$0.25/page

Project Stats

- Number of pages per year
 - POC: 1M pages
 - PROD: 5M pages
- Number of properties
 - POC: 1,800
 - PROD: 8,300

Use Case: Classify Out Pages

150 page documents but only 20 pages need to be extracted

- 150 pages x \$0.03/page is \$4.50
- OR
- 150 pages x \$0.003/page for classification at \$0.45
- 20 pages x \$0.03/page for extraction equals \$0.60
- RESULTS
- 76% savings with this approach
- \$4.50 vs. \$1.05



Azure OpenAI Generative AI for Medical Charts

Use Case #5: Vital Chart Gen AI QC Process

The Customer

- VRC performs medical scanning and coding services using manual labor

The Challenge

- Ensure that patient charts have ONLY that patient documents, no other patients
- Provider documentation is completely unstructured
- Building models would take too long and require lots of PHI samples

The Solution

- Query prompt engineering
- Kodak Info Input capture wrapper
- Azure OpenAI



Why Azure OpenAI Over AWS?

- Azure did a better job identifying Dates of Service in the context of NLP





Live Demonstrations