



CDW Documentation

CLI ML Workspace Transcript

CLI ML Workspace Transcript

In **Azure Machine Learning (Azure ML)**, datasets used for training can be in a variety of formats. The platform is flexible, but depending on how you plan to use the dataset (e.g., for AutoML, custom training, or pipelines), certain formats are more appropriate.

□ Commonly Supported Formats

Format	Description	Notes
CSV (.csv)	Comma-separated values	Widely used for tabular data; required for AutoML classification/regression tasks
TSV (.tsv)	Tab-separated values	Same as CSV but tab-delimited
Parquet (.parquet)	Columnar storage optimized for big data	Efficient for large-scale tabular data
JSON (.json)	JavaScript Object Notation	Used for hierarchical or semi-structured data
Image folders	Directory of images	For vision tasks; Azure ML expects structure like <code>dataset/class_name/image.jpg</code>
TFRecord (.tfrecord)	TensorFlow Record	Used for TensorFlow models
Pickle (.pkl)	Python serialized objects	Only for trusted pipelines; not suitable for AutoML
DataFrames	Pandas/Spark dataframes	When passing data in memory in notebooks or pipelines

□ Best Practices for Training in Azure ML

Scenario	Recommended Format
AutoML Tabular	CSV or Parquet
Computer Vision	Folder-based image dataset or JSON with image metadata
NLP (Text)	CSV, TSV, or JSON with text columns
Custom ML Training (PyTorch, TensorFlow, etc.)	Any format you can load in your script, usually CSV, Parquet, or TFRecord
Time Series Forecasting	CSV with datetime and target columns

□ Where to Store

Datasets must be accessible to Azure ML compute. Typical storage options:

- Azure Blob Storage (via Datastore)
- Azure Files
- Local uploads (less common, mostly for small experiments)
- Public web URLs

□ Dataset Registration (Optional but Useful)

You can register your dataset in Azure ML for reuse:

```
pythonCopyEditfrom azureml.core import Workspace, Dataset

ws = Workspace.from_config()
datastore = ws.get_default_datastore()

dataset = Dataset.Tabular.from_delimited_files(path=(datastore,
'path/to/data.csv'))
dataset = dataset.register(workspace=ws, name='my_dataset',
create_new_version=True)
```

Would you like an example setup for a specific task (e.g., AutoML image classification, time series forecasting, or PyTorch model training)?

You said:

mysql command to dump database to csv file?

ChatGPT said:

MySQL does **not natively export to CSV** directly via mysqldump. However, you can export tables to CSV using SQL queries or command-line tools like mysql + INTO OUTFILE.

□ Option 1: Export a Table to CSV Using SQL (INTO OUTFILE)

```
sqlCopyEditSELECT *
INTO OUTFILE '/var/lib/mysql-files/my_table.csv'
FIELDS TERMINATED BY ','
ENCLOSED BY '"'
LINES TERMINATED BY '\n'
FROM my_table;
```

Notes:

- The path **must be writable by MySQL** and usually inside /var/lib/mysql-files/.
- You must have **FILE privileges**.
- This method works **per table**, not the entire database in one shot.

□ Option 2: Use mysql CLI and sed/awk (CSV-like output)

```
bashCopyEditmysql -u username -p -e "SELECT * FROM my_database.my_table;" \
```

```
| sed 's/\t/", "/g;s/^"/";s/$"/' > my_table.csv
```

This simulates CSV by converting tab-separated output to quoted CSV.

□ Option 3: Python Script to Dump Entire Database to CSV Files (per Table)

If you want a full database export (one CSV per table):

```
pythonCopyEditimport mysql.connector
import pandas as pd

conn = mysql.connector.connect(
    host="localhost", user="your_user", password="your_pass",
    database="your_db"
)

cursor = conn.cursor()
cursor.execute("SHOW TABLES;")
tables = cursor.fetchall()

for (table_name,) in tables:
    df = pd.read_sql(f"SELECT * FROM {table_name}", conn)
    df.to_csv(f"{table_name}.csv", index=False)

conn.close()
```

□ Permissions Reminder

If you're getting errors with INTO outfile, check:

- The **secure_file_priv** system variable:

```
sqlCopyEditSHOW VARIABLES LIKE 'secure_file_priv';
```

- The MySQL user has FILE permission:

```
sqlCopyEditGRANT FILE ON *.* TO 'your_user'@'localhost';
```

Would you like to export all tables in a DB to a **single CSV**, or do you prefer **one file per table**?

You said:

how to upload dataset to Azure ML?

ChatGPT said:

To **upload a dataset to Azure Machine Learning (Azure ML)**, you have multiple options depending on your workflow. Here are the most common methods:

□ 1. Via Azure ML Studio (Web UI)

□ Steps:

1. Go to Azure ML Studio
2. Navigate to:
Data → + Create → From local files
3. Fill in:
 - **Name** and **Description**
 - **Type**: *Tabular* or *File*
 - **Browse and upload file(s)** (e.g., CSV, image folders)
 - Choose a **datastore** (e.g., workspaceblobstore)
 - Click **Next**, configure schema (if tabular), then click **Create**

Best for one-time manual uploads of CSVs, images, etc.

□ 2. Via Python SDK

□ Example: Upload a CSV as a tabular dataset

```
pythonCopyEditfrom azureml.core import Workspace, Dataset, Datastore

ws = Workspace.from_config()
datastore = ws.get_default_datastore()

# Upload file to datastore
datastore.upload_files(
    files=['./local_data/mydata.csv'],
    target_path='datasets/mydata/',
    overwrite=True,
)

# Create Tabular Dataset
dataset = Dataset.Tabular.from_delimited_files(path=(datastore,
'datasets/mydata/mydata.csv'))

# Optional: Register dataset for reuse
dataset.register(
    workspace=ws,
```

```

name='my_csv_dataset',
description='Uploaded via SDK',
create_new_version=True
)

```

Best for scripting and versioning datasets for pipelines or experiments.

□ 3. Via Azure CLI

```

bashCopyEditaz ml data create --name my_csv_dataset \
--type uri_file \
--path ./local_data/mydata.csv \
--workspace-name my-ml-workspace \
--resource-group my-rg

```

Good for automation and integration with CI/CD tools.

□ 4. From Azure Blob Storage or External URL

If your data is already in Azure Blob:

```

pythonCopyEdit# Reference without uploading
from azureml.core import Dataset

datastore = ws.get_default_datastore()

dataset = Dataset.Tabular.from_delimited_files(path=(datastore,
'path/in/blob/file.csv'))
dataset.register(workspace=ws, name="external_dataset",
create_new_version=True)

```

Or from a public URL:

```

pythonCopyEditdataset =
Dataset.Tabular.from_delimited_files("https://example.com/data.csv")

```

□ Upload Types

Type	Use Case
Tabular	CSV, TSV, Parquet - for structured data
File	Images, binaries, JSON, PDFs - any file type

Would you like an example tailored for **AutoML**, **Vision**, or **training with your own script**?

You said:

collectiblesNow interactive!

id	character	figure_name	property
1	Aang	Nendroid with Momo	Avatar the Last Airbender
2	Alice	D-Stage Alice	Alice In Wonderland
3	Alice	Disney Animators Collection Alice Doll	Alice In Wonderland
4	Alice	Disney Store Classic	Alice In Wonderland

Given this data, how would configure and train ML to let me query against a character and get the resulting details from the dataset?

ChatGPT said:

Great — you've uploaded `collectibles.csv`. To **configure and train a machine learning model to query character details**, it sounds like you're aiming for **intelligent search or semantic lookup**, rather than traditional predictive modeling.

Here's a breakdown of how to do this in Azure ML:

□ Goal

You want to **input a character name** and retrieve **associated details** from your dataset. This suggests:

- Text-based lookup or semantic matching
- Possibly embedding-based retrieval (if character names aren't exact matches)

□ Step-by-Step Plan (Using Azure ML + OpenAI or Embeddings)

Step 1: Understand Your Data

I'll examine your dataset.

Your dataset contains detailed collectibles information with fields like:

- character

- figure_name
 - property
 - type
 - manufacturer
 - location
 - quantity, list_price, total_cost, approximate_value
-

□ ML-Based Query System Options

Since you're aiming to **query by character** and get full details, you have two main paths:

□ Option 1: Traditional Filter-Based Query

No ML needed — just build a UI or API to filter rows where `character = input`.

Best for exact or dropdown-based search.

□ Option 2: AI-Powered Search (Semantic or Fuzzy Matching)

Use **OpenAI embeddings** + **Azure ML** to let the user type:

```
"Find me figures from Alice in Wonderland" or "What do you have of Aang?"
```

Steps:

Step 1: Generate Embeddings for Each Row

Use OpenAI's `text-embedding-ada-002` or Azure OpenAI equivalent on a string like:

```
pythonCopyEditrow_text = f"{character}, {figure_name}, {property}, {type}, {manufacturer}"
```

Step 2: Store Embeddings in a vector index (e.g., FAISS or Azure Cognitive Search)

Step 3: User types a query → convert to embedding → run similarity search → return top results.

□ Recommendation

Would you prefer:

1. **Exact match web app / CLI tool** (quick lookup, no ML)?
2. **AI-powered search** using semantic embeddings?
3. **Train a basic classification/regression model** (less suitable here)?

I can generate code or deploy scripts for either path. Which would you like?

You said:

I would like to use ML for this. Do you have suggestions for some example that would be more suited to using ML with this dataset?

ChatGPT said:

Great — your dataset is rich and could support several ML use cases. While direct character lookup is more of a search task (good for semantic matching), we can design **ML training tasks** that **learn patterns** from the collectibles.

Here are the most suitable ML use cases for your dataset:

□ 1. Price Estimation (Regression)

□ **Problem:**

Predict `approximate_value` based on other fields (`character`, `type`, `manufacturer`, `list_price`, etc.).

□ **Why it fits:**

- You have labeled numeric data.
- It's useful for appraising unlisted items.

□ **ML Type:**

- **Regression** (e.g., Linear Regression, XGBoost, AutoML)
-

□ 2. Category Prediction (Classification)

□ Problem:

Predict the type of collectible or its location based on fields like character, property, and manufacturer.

□ Why it fits:

- Helps auto-label new entries.
- Enables predictive organization/sorting.

□ ML Type:

- **Classification**
-

□ 3. Clustering Characters by Similarity (Unsupervised)

□ Problem:

Group similar characters or figures automatically to find patterns or suggest groupings.

□ Why it fits:

- Explores hidden structure.
- Could be visualized or used to build recommendation systems.

□ ML Type:

- **Clustering (e.g., KMeans)** using embeddings of text descriptions.
-

□ 4. Anomaly Detection

□ Problem:

Find collectibles that are priced much higher/lower than similar items.

□ Why it fits:

- Helps find mispriced entries or rare items.

□ ML Type:

- **Anomaly Detection (Isolation Forest, One-Class SVM)**
-

□ 5. Embedding-Based Semantic Search (LLM-enhanced ML)

□ Problem:

Let users search using natural language and get intelligent matches.

□ Why it fits:

- Transforms character/item descriptions into embeddings.
- Allows vector similarity search over the dataset.

□ Tools:

- Azure OpenAI (for embeddings)
 - FAISS or Azure Cognitive Search (for similarity search)
 - Optional: fine-tuning a retrieval model
-

□ Suggestion

To get started with a real ML training task:

► Start with Price Estimation using Regression:

- **Input:** character, figure_name, property, type, manufacturer, list_price, etc.
- **Output:** approximate_value

Would you like me to generate a full training script (e.g., in Python with scikit-learn or in Azure ML SDK) for this?

You said:

Yes

ChatGPT said:

Your regression model has been successfully trained to predict `approximate_value` using collectible attributes. The evaluation shows:

□ **Root Mean Squared Error (RMSE): ~30.15**

This means that, on average, the model's predictions are off by about **\$30** — not bad for a first pass, considering potential variance in collectible values.

□ **Next Steps You Can Take**

□ **Improvements**

- Use more advanced models (e.g., **XGBoost**, **LightGBM**).
- Add more features (e.g., `quantity`, `location`, or textual embeddings).
- Clean or normalize text fields (e.g., standardize "Alice in Wonderland").

□ **Use in Azure ML**

To deploy this in **Azure Machine Learning**, you can:

- Package the training in a script (`train.py`)
- Use an Azure ML pipeline or AutoML for automated tuning
- Deploy the trained model as a **REST endpoint** for use in apps

Would you like a downloadable Python script (`train.py`) to run this on your machine or in Azure ML Notebook/Compute?

You said:

yes, `train.py` please

ChatGPT said:

□ Your `train.py` script is ready for download:

□ [Click here to download train.py](#)

This script:

- Loads `collectibles.csv`
- Trains a regression model to predict `approximate_value`

- Prints the RMSE
- Saves the model as `collectibles_model.pkl`

Let me know if you'd like a deployment script, inference example, or Azure ML pipeline wrapper for it.

You said:

deployment script and inference exmple

ChatGPT said:

Here are your deployment and inference scripts:

- `Download deploy.py` — Launches a FastAPI server that loads the trained model and exposes a `/predict` endpoint.
- `Download test_inference.py` — Sends a test prediction request to the local API.

📄 How to Use

1. Start the API Server

```
bashCopyEditpip install fastapi uvicorn joblib scikit-learn pandas
uvicorn deploy:app --reload
```

2. Run the Inference Example

```
bashCopyEditpython test_inference.py
```

It should return:

```
jsonCopyEdit{"approximate_value": <predicted_value>}
```

Would you like an Azure ML deployment version (e.g., scoring script + inference config for an online endpoint)?

You said:

Yes. I would like an Azure ML Version

ChatGPT said:

Here are the files you need for deploying your model to **Azure Machine Learning**:

□ Download Deployment Files

- □ score.py — Scoring script used by the inference endpoint.
- ⚙ inference_config.json — Configuration for the deployment.
- □ env.yml — Conda environment definition.

□ How to Deploy in Azure ML (Python SDK v1 or v2)

Here's a simplified v2-style Azure ML deployment flow:

```
pythonCopyEditfrom azure.ai.ml import MLClient
from azure.ai.ml.entities import ManagedOnlineEndpoint,
ManagedOnlineDeployment
from azure.identity import DefaultAzureCredential
from azure.ai.ml.entities import Environment, Model

ml_client = MLClient(DefaultAzureCredential(), "SUBSCRIPTION_ID",
"RESOURCE_GROUP", "WORKSPACE")

# Register the model
model = ml_client.models.create_or_update(
    Model(name="collectibles-model", path="collectibles_model.pkl")
)

# Create environment
env = Environment(
    name="collectibles-env",
    conda_file="env.yml",
    description="Env for collectibles model"
)
ml_client.environments.create_or_update(env)

# Create endpoint
endpoint = ManagedOnlineEndpoint(
    name="collectibles-endpoint",
    auth_mode="key"
)
ml_client.begin_create_or_update(endpoint).result()

# Create deployment
deployment = ManagedOnlineDeployment(
    name="blue",
    endpoint_name=endpoint.name,
    model=model,
    environment=env,
    code_configuration={"code": ".", "scoring_script": "score.py"},
    instance_type="Standard_DS2_v2",
```

```
    instance_count=1
)
ml_client.begin_create_or_update(deployment).result()

# Set deployment as default
ml_client.online_endpoints.invoke(
    endpoint_name="collectibles-endpoint",
    request_file="test_sample.json"
)
```

Would you like a zipped bundle with all files needed to deploy this in Azure ML Studio or CLI?

You said:

yes

ChatGPT said:

The error indicates that the trained model file `collectibles_model.pkl` is missing. Let me regenerate and save the model file first, then I'll recreate the deployment bundle.

☐ Your Azure ML deployment bundle is ready:

☐ Download `azure_ml_deployment_bundle.zip`

This ZIP file contains:

- `score.py` - Model inference logic
- `env.yml` - Conda environment for deployment
- `inference_config.json` - Metadata for Azure deployment
- `collectibles_model.pkl` - Trained model

Let me know if you'd like a YAML-based deployment template or SDK script to automate this deployment on Azure ML.

You said:

sdk script and instructions on how to use it

ChatGPT said:

☐ Your Azure ML SDK deployment script is ready:

☐ Download `deploy_to_azureml.py`

□ How to Use This Script

□ Prerequisites

1. Install Azure ML SDK

```
bashCopyEditpip install azure-ai-ml azure-identity
```

2. Set up Azure CLI & Login

```
bashCopyEditaz login
```

3. Set the following in the script:

- `subscription_id` = "your-subscription-id"
- `resource_group` = "your-resource-group"
- `workspace_name` = "your-ml-workspace"

4. Place the following files in the same folder:

- `collectibles_model.pkl`
- `score.py`
- `env.yml`

▶ Run the Script

```
bashCopyEditpython deploy_to_azureml.py
```

□ To Query the Model After Deployment

Here's how you can query the endpoint once it's live:

```
pythonCopyEditfrom azure.ai.ml import MLClient
from azure.identity import DefaultAzureCredential

ml_client = MLClient(
    DefaultAzureCredential(),
    "your-subscription-id",
    "your-resource-group",
    "your-ml-workspace"
)

response = ml_client.online_endpoints.invoke(
    endpoint_name="collectibles-endpoint",
    request_file="test_sample.json" # JSON list of dicts
)
```

```
print(response)
```

Would you like a ready-made test_sample.json to use with this?

You said:

Got an error during deployment: Traceback (most recent call last): File
"/Users/don.dehamer/Azure/./deploy_to_azureml.py", line 33, in <module> registered_model =
ml_client.models.create_or_update(model) File
"/Users/don.dehamer/.local/pipx/venvs/requests/lib/python3.13/site-
packages/azure/ai/ml/_telemetry/activity.py", line 288, in wrapper return f(*args, kwargs) **File**
"/Users/don.dehamer/.local/pipx/venvs/requests/lib/python3.13/site-
packages/azure/ai/ml/operations/_model_operations.py", line 301, in create_or_update
raise ex File "/Users/don.dehamer/.local/pipx/venvs/requests/lib/python3.13/site-
packages/azure/ai/ml/operations/_model_operations.py", line 255, in create_or_update
validate_ml_flow_folder(model.path, model.type) # type: ignore
~~~~~**File**  
"/Users/don.dehamer/.local/pipx/venvs/requests/lib/python3.13/site-  
packages/azure/ai/ml/\_utils/utils.py", line 569, in validate\_ml\_flow\_folder raise  
ValidationException( ...<4 lines>... ) azure.ai.ml.exceptions.ValidationException: Error with  
path azureml:subscriptions/baa29726-b3e6-4910-bb9b-  
b585c655322c/resourcegroups/don-test-rg-SCUS/workspaces/don-ml-workspace-  
new/datstores/workspaceblobstore/paths/LocalUpload/6eb9cb11e8d540eebeea23fdd3b3e  
8b0/collectibles\_model.pkl. Model of type mlflow\_model cannot register a file. ChatGPT  
said: Thanks for the detailed error — this happens because the model was registered as  
mlflow\_model but the artifact (collectibles\_model.pkl) is just a regular pickle file, not  
an MLflow directory. -- =====  Fix: Register it as a custom\_model instead =====  
Update this section in your deploy\_to\_azureml.py script: =====  Before: ===== <code →  
pythonCopyEditmodel = Model( path="collectibles\_model.pkl", name="collectibles-  
model", description="Predicts approximate value of a collectible", type="mlflow\_model" )  
</code> =====  After: ===== <code → pythonCopyEditmodel = Model(  
path="collectibles\_model.pkl", name="collectibles-model", description="Predicts  
approximate value of a collectible", type="custom\_model" # Correct type for a pickle file )  
</code> -- =====  Why this matters: ===== \* mlflow\_model expects a full MLflow  
directory with MLmodel, conda.yaml, etc. \* custom\_model is correct for plain .pkl or  
joblib-based models. -- Would you like me to regenerate the corrected script and upload  
it again for convenience? == You said: == New error: File  
"/Users/don.dehamer/Azure/./deploy\_to\_azureml.py", line 42, in <module>  
ml\_client.environments.create\_or\_update(env)  
~~~~~**File**  
"/Users/don.dehamer/.local/pipx/venvs/requests/lib/python3.13/site-
packages/azure/ai/ml/_telemetry/activity.py", line 288, in wrapper return f(*args, kwargs)
File "/Users/don.dehamer/.local/pipx/venvs/requests/lib/python3.13/site-
packages/azure/ai/ml/operations/_environment_operations.py", line 212, in
create_or_update raise ex File
"/Users/don.dehamer/.local/pipx/venvs/requests/lib/python3.13/site-
packages/azure/ai/ml/operations/_environment_operations.py", line 123, in
create_or_update return self._try_update_latest_version(next_version, latest_version,
environment)

ACR for storing environments and containers. This ACR: * Was either deleted manually, or * The workspace was misconfigured or corrupted --- ===== **How to Fix** =====

Option 1: Check Workspace Configuration ===== - **Go to the Azure Portal - Open your ML workspace: don-ml-workspace-new - Under "Linked resources", verify: * A Container Registry is listed * Its name resolves properly and hasn't been deleted** If it's missing, the workspace is broken and cannot run deployments. --- ===== **Option 2: Create a New ML Workspace** ===== If the ACR is missing or deleted and cannot be restored: - **In Azure Portal: * Go to ML → Create new workspace * Name it something like don-ml-new * Choose region (e.g., South Central US) - Let Azure auto-create all resources (Storage, ACR, Key Vault, App Insights) Then in your script, update:** `<code> pythonCopyEditworkspace_name = "don-ml-new" resource_group = "don-test-rg-SCUS" </code>` --- ===== **Option 3: Manually Link a New ACR** ===== **Advanced, not recommended unless you're reconfiguring existing infra: * Create a new ACR * Link it to the ML workspace using ARM or CLI** --- ===== **Recommended Next Step** ===== **Do you want: * A new workspace creation script * Or just continue with a manually created ML workspace and retry? Let me know and I can walk you through it or update the code accordingly. == You said: == I don't see a linked resources option in the menu on the left. ChatGPT said: Thanks for checking — Microsoft has recently reorganized the Azure ML workspace interface, and "Linked resources" may not appear as a standalone menu item anymore, especially in some regions or subscription tiers. But you can still verify if the Azure ML workspace has a valid ACR by using one of these methods: --- ===== Method 1: Check in Azure Portal via "Deployment Center" ===== - Navigate to your Azure ML Workspace (don-ml-workspace-new) - In the left-hand menu, open: * Deployment Center or Endpoints → Attempt to deploy manually - If the workspace is missing an ACR, you'll get a visible prompt or error message during this process. --- ===== Method 2: Use Azure CLI to Inspect Linked ACR ===== Run this command: `<code> bashCopyEditaz ml workspace show \ -name don-ml-workspace-new \ -resource-group don-test-rg-SCUS \ -query container_registry </code>` This should return something like: `<code> jsonCopyEdit"/subscriptions/<sub-id>/resourceGroups/<rg>/providers/Microsoft.ContainerRegistry/registries/<acr-name>" </code>` If it returns null or a broken ACR ID, the workspace is misconfigured. --- ===== Method 3: Use Azure Resource Graph to Search ===== You can also query your subscription's ACR resources: `<code> bashCopyEditaz acr list -resource-group don-test-rg-SCUS -output table </code>` Make sure there's an ACR available and not deleted. --- ===== What to Do if ACR is Missing ===== If your workspace's ACR is missing or invalid: ===== **Best Fix: Recreate the Workspace** ===== The simplest and most reliable fix is to create a new ML workspace, letting Azure create a new ACR with it automatically. You can do that in the portal or with: `<code> bashCopyEditaz ml workspace create \ -name don-ml-workspace-fixed \ -resource-group don-test-rg-SCUS \ -location southcentralus </code>` --- **Would you like me to provide: * A CLI script to recreate the ML workspace? * An updated SDK deploy script with validation for ACR presence before deployment? Let me know how you want to proceed. == You said: == I got: az ml workspace show \ -name don-ml-workspace-new \ -resource-group don-test-rg-SCUS \ -query container_registry "/subscriptions/baa29726-b3e6-4910-bb9b-b585c655322c/resourceGroups/don-test-rg-SCUS/providers/Microsoft.ContainerRegistry/registries/1bc3bfe4a5644d918283c13025dc9d8c" ChatGPT said: Perfect — that confirms your ML workspace is still pointing to the ACR with ID: `<code> bashCopyEditMicrosoft.ContainerRegistry/registries/1bc3bfe4a5644d918283c13025dc9d8c </code>` But this ACR appears to be either: * Deleted, or * Inaccessible (e.g., region mismatch, permission issue, or soft-deleted) --- ===== Next Steps to Resolve =====****

==== **Step 1: Check if ACR Exists** ==== **Run:** `bashCopyEditaz acr show \ -ids "/subscriptions/baa29726-b3e6-4910-bb9b-b585c655322c/resourceGroups/don-test-rg-SCUS/providers/Microsoft.ContainerRegistry/registries/1bc3bfe4a5644d918283c13025dc9d8c"` **====** **If you get an error:** `>ResourceNotFound` — the ACR is deleted and Azure ML won't work until this is fixed. -- **====** **Step 2: Recreate the Workspace with a Fresh ACR (Recommended Fix)** **====** `bashCopyEditaz ml workspace create \ -name don-ml-workspace-fixed \ -resource-group don-test-rg-SCUS \ -location southcentralus` Azure will automatically: * Create a valid ACR * Link it to the workspace * Allow deployments to proceed -- **====** **Then Update Your Python Deployment Script** **====** **Change this section:** `pythonCopyEditworkspace_name = "don-ml-workspace-fixed" # Updated name resource_group = "don-test-rg-SCUS"` Everything else remains the same. -- **====** **Alternative (Advanced): Link a New ACR Manually** **====** **If you want to avoid recreating the workspace (not recommended unless necessary), you can:** - Create a new ACR:`bashCopyEditaz acr create -name mynewacr123 -resource-group don-test-rg-SCUS -sku Basic` - Manually patch the ML workspace to use it (involves ARM template or REST call — fragile). -- **Would you like:** * A CLI script to recreate the workspace and validate it? * Or an updated deployment script ready for the new workspace? **Let me know your preference.** == **You said:** == `az acr show \ -ids "/subscriptions/baa29726-b3e6-4910-bb9b-b585c655322c/resourceGroups/don-test-rg-SCUS/providers/Microsoft.ContainerRegistry/registries/1bc3bfe4a5644d918283c13025dc9d8c"` the following arguments are required: `-name/-n` **ChatGPT said:** Ah — that error means `az acr show` does not support the `-ids` parameter alone in this context, despite `az resource show` supporting it. Let's fix that by extracting and using the ACR name directly. -- **====** **Step-by-Step Fix** **====** **====** **Step 1: Extract ACR Name from the full ID** **====** **From your previous output:** `bashCopyEdit.../registries/1bc3bfe4a5644d918283c13025dc9d8c` **So the ACR name is:** `bashCopyEdit1bc3bfe4a5644d918283c13025dc9d8c` -- **====** **Step 2: Run This to Check If the ACR Exists** **====** `bashCopyEditaz acr show \ -name 1bc3bfe4a5644d918283c13025dc9d8c \ -resource-group don-test-rg-SCUS` -- **====** **Expected Outcome** **====** * If the ACR exists: You'll get a JSON with details like location, login server, etc. * If it doesn't: You'll get:`bashCopyEdit(ResourceNotFound) The Resource 'Microsoft.ContainerRegistry/registries/1bc3bfe4a5644d918283c13025dc9d8c' under resource group 'don-test-rg-SCUS' was not found.` Let me know what it returns — that will confirm the next step: reuse the workspace or recreate it. == **You said:** == `az acr show \ -name 1bc3bfe4a5644d918283c13025dc9d8c \ -resource-group don-test-rg-SCUS` **(ResourceNotFound) The Resource 'Microsoft.ContainerRegistry/registries/1bc3bfe4a5644d918283c13025dc9d8c' under resource group 'don-test-rg-SCUS' was not found. For more details please go to <https://aka.ms/ARMResourceNotFoundFix> Code: ResourceNotFound Message: The Resource 'Microsoft.ContainerRegistry/registries/1bc3bfe4a5644d918283c13025dc9d8c' under resource group 'don-test-rg-SCUS' was not found. For more details please go to <https://aka.ms/ARMResourceNotFoundFix>** **ChatGPT said:** Thanks — this confirms that the Azure Container Registry (ACR) linked to your ML workspace has been deleted or is permanently unavailable, and as a result, your current workspace cannot deploy endpoints or environments. -- **====** **Final Fix: Create a New Workspace (Auto-links a Valid ACR)** **====** **====** **Step 1: Create a New Azure ML Workspace** **====** **Run the following command:** `bashCopyEditaz ml workspace create \ -name don-ml-workspace-fixed \ -resource-group don-test-rg-SCUS \ -location southcentralus` **>This will automatically create and link a working ACR, storage, key vault, and app**

insights instance to the new ML workspace. -- ===== **Step 2: Update Your Deployment Script =====** Update these two lines in `deploy_to_azureml.py`: `<code> pythonCopyEditworkspace_name = "don-ml-workspace-fixed" resource_group = "don-test-rg-SCUS" </code>` Then re-run your script: `<code> bashCopyEditpython deploy_to_azureml.py </code>` **-- =====** **After Deployment =====** Once it succeeds, you can: *** Use `ml_client.online_endpoints.invoke(...)` to query it * Manage the endpoint in Azure ML Studio** **-- Would you like me to repackage the deployment script and endpoint test script with the new workspace values set for convenience? == You said: ==** **New error with new environment: Traceback (most recent call last): File "/Users/don.dehamer/Azure/./deploy_to_azureml.py", line 52, in <module> ml_client.begin_create_or_update(endpoint).result()**

```

~~~~~
^^^^^^^^^^^^^^^^ File
"/Users/don.dehamer/.local/pipx/venvs/requests/lib/python3.13/site-
packages/azure/ai/ml/_ml_client.py", line 1308, in begin_create_or_update return
_begin_create_or_update(entity, self._operation_container.all_operations, kwargs) File
"/opt/homebrew/Cellar/python@3.13/3.13.3/Frameworks/Python.framework/Versions/3.13/l
ib/python3.13/functools.py", line 934, in wrapper return dispatch(args[0].class)(*args, kw)
~~~~~
^^^^^^^^^^^^^^^^ File
"/Users/don.dehamer/.local/pipx/venvs/requests/lib/python3.13/site-
packages/azure/ai/ml/_ml_client.py", line 1408, in _return
operations[AzureMLResourceType.ONLINE_ENDPOINT].begin_create_or_update(entity,
kwargs)
~~~~~
^^^^^^^^^^^^^^^^ File
"/Users/don.dehamer/.local/pipx/venvs/requests/lib/python3.13/site-
packages/azure/core/tracing/decorator.py", line 119, in wrapper_use_tracer return
func(*args, kwargs) File
"/Users/don.dehamer/.local/pipx/venvs/requests/lib/python3.13/site-
packages/azure/ai/ml/_telemetry/activity.py", line 288, in wrapper return f(*args, kwargs)
File "/Users/don.dehamer/.local/pipx/venvs/requests/lib/python3.13/site-
packages/azure/ai/ml/operations/_online_endpoint_operations.py", line 263, in
begin_create_or_update raise ex File
"/Users/don.dehamer/.local/pipx/venvs/requests/lib/python3.13/site-
packages/azure/ai/ml/operations/_online_endpoint_operations.py", line 258, in
begin_create_or_update raise ex File
"/Users/don.dehamer/.local/pipx/venvs/requests/lib/python3.13/site-
packages/azure/ai/ml/operations/_online_endpoint_operations.py", line 247, in
begin_create_or_update poller = self._online_operation.begin_create_or_update(
resource_group_name=self._resource_group_name, ...<4 lines>... self._init_kwargs, ) File
"/Users/don.dehamer/.local/pipx/venvs/requests/lib/python3.13/site-
packages/azure/core/tracing/decorator.py", line 119, in wrapper_use_tracer return
func(*args, kwargs) File
"/Users/don.dehamer/.local/pipx/venvs/requests/lib/python3.13/site-
packages/azure/ai/ml/_restclient/v2022_02_01_preview/operations/_online_endpoints_opera
tions.py", line 921, in begin_create_or_update raw_result = self._create_or_update_initial(
resource_group_name=resource_group_name, ...<5 lines>... kwargs ) File
"/Users/don.dehamer/.local/pipx/venvs/requests/lib/python3.13/site-
packages/azure/ai/ml/_restclient/v2022_02_01_preview/operations/_online_endpoints_opera
tions.py", line 856, in _create_or_update_initial raise
HttpResponseError(response=response, error_format=ARMErrrorFormat)
azure.core.exceptions.HttpResponseError: (BadRequest) The request is invalid. Code:

```



```

packages/azure/core/polling/base_polling.py", line 788, in run self._poll()
~~~~~^^ File "/Users/don.dehamer/.local/pipx/venvs/requests/lib/python3.13/site-
packages/azure/core/polling/base_polling.py", line 820, in _poll raise
OperationFailed("Operation failed or canceled")
azure.core.polling.base_polling.OperationFailed: Operation failed or canceled The above
exception was the direct cause of the following exception: Traceback (most recent call
last): File "/Users/don.dehamer/Azure/./deploy_to_azureml.py", line 67, in <module>
ml_client.begin_create_or_update(deployment).result()
~~~~~^^ File
"/Users/don.dehamer/.local/pipx/venvs/requests/lib/python3.13/site-
packages/azure/core/polling/_poller.py", line 254, in result self.wait(timeout)
~~~~~^^^^^^^^ File
"/Users/don.dehamer/.local/pipx/venvs/requests/lib/python3.13/site-
packages/azure/core/tracing/decorator.py", line 119, in wrapper_use_tracer return
func(*args, kwargs) File
"/Users/don.dehamer/.local/pipx/venvs/requests/lib/python3.13/site-
packages/azure/core/polling/_poller.py", line 273, in wait raise self._exception # type:
ignore ^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^ File
"/Users/don.dehamer/.local/pipx/venvs/requests/lib/python3.13/site-
packages/azure/core/polling/_poller.py", line 188, in _start self._polling_method.run()
~~~~~^^ File
"/Users/don.dehamer/.local/pipx/venvs/requests/lib/python3.13/site-
packages/azure/core/polling/base_polling.py", line 803, in run raise
HttpResponseError(response=self._pipeline_response.http_response, error=err) from err
azure.core.exceptions.HttpResponseError: (ResourceNotReady) User container has
crashed or terminated. Please see troubleshooting guide, available here:
https://aka.ms/oe-tsg#error-resourcenotready Code: ResourceNotReady Message: User
container has crashed or terminated. Please see troubleshooting guide, available here:
https://aka.ms/oe-tsg#error-resourcenotready ChatGPT said: You're almost there — this
final error: >(ResourceNotReady) User container has crashed or terminated means
the endpoint deployed, but your scoring container crashed at runtime. -- ===== □ Root
Cause (Most Likely) ===== Your score.py script depends on packages like pandas,
scikit-learn, or joblib, but: * The base image + environment doesn't install them
correctly, or * The env.yml has a compatibility issue. -- ===== □ Step-by-Step Fix =====
===== □ Step 1: Use a Known-Good Azure ML Base Image ===== Update the Environment
definition to use one of Azure's prebuilt, fully compatible images: <code →
pythonCopyEditimage="mcr.microsoft.com/azureml/openmpi4.1.0-ubuntu20.04" </code>
You're likely already doing this — □ good. -- ===== □ Step 2: Verify/Update Your env.yml
===== Ensure it's minimal but complete: <code → yamlCopyEditname: collectibles-env
dependencies: - python=3.8 - pip - pip: - pandas - scikit-learn - joblib </code> >□ If you're
using Python 3.13 locally, be aware Azure ML does not support 3.13 yet. Stick to 3.8 or 3.9
for compatibility. -- ===== □ Step 3: Add Logging to score.py ===== Wrap your code in
try/except and add logging so you can see container logs via the Azure ML Studio portal:
<code → pythonCopyEditimport joblib import pandas as pd import json import os import
logging def init(): global model try: model_path =
os.path.join(os.getenv('AZUREML_MODEL_DIR'), 'collectibles_model.pkl') model =
joblib.load(model_path) except Exception as e: logging.exception("Model initialization
failed") def run(raw_data): try: data = pd.DataFrame(json.loads(raw_data)) predictions =
model.predict(data) return predictions.tolist() except Exception as e:
logging.exception("Inference failed") return str(e) </code> -- ===== □ Step 4: Redeploy
===== - Update the env.yml and score.py - Repackage or re-run the deployment - If it

```


show errors in logs and never crash the container silently: `→ pythonCopyEditimport joblib import pandas as pd import json import os import logging def init(): global model try: model_path = os.path.join(os.getenv("AZUREML_MODEL_DIR"), "collectibles_model.pkl") model = joblib.load(model_path) logging.info(" Model loaded successfully.") except Exception as e: logging.exception(" Model loading failed") def run(raw_data): try: logging.info(f"Received input: {raw_data}") data = pd.DataFrame(json.loads(raw_data)) predictions = model.predict(data) return {"predictions": predictions.tolist()} except Exception as e: logging.exception(" Inference failed") return {"error": str(e)} </code> -- ===== Double-Check: Final env.yml =====`

<code → ymlCopyEditname: collectibles-env dependencies: - python=3.9 - pip - pip: - pandas - scikit-learn - joblib </code> -- ===== If Logs Say "Missing module" or "Cannot import..." ===== Then we're back to missing dependencies. Let me know and I'll regenerate an updated version of: * env.yml * score.py * A new ZIP bundle * A redeployment script -- Just let me know what's in the Azure ML container logs, and we'll squash this for good. == You said: == 2025-06-03T17:16:57,809228388+00:00 | gunicorn/run | 2025-06-03T17:16:57,810899924+00:00 | gunicorn/run | Entry script directory: /var/azureml-app/Azure/. 2025-06-03T17:16:57,812821682+00:00 | gunicorn/run | 2025-06-03T17:16:57,814621429+00:00 | gunicorn/run | ##### 2025-06-03T17:16:57,816265263+00:00 | gunicorn/run | Dynamic Python Package Installation 2025-06-03T17:16:57,817917798+00:00 | gunicorn/run | ##### 2025-06-03T17:16:57,819484327+00:00 | gunicorn/run | 2025-06-03T17:16:57,821439186+00:00 | gunicorn/run | Dynamic Python package installation is disabled. 2025-06-03T17:16:57,823101522+00:00 | gunicorn/run | 2025-06-03T17:16:57,824772959+00:00 | gunicorn/run | ##### 2025-06-03T17:16:57,826426494+00:00 | gunicorn/run | Checking if the Python package azureml-inference-server-http is installed 2025-06-03T17:16:57,828358152+00:00 | gunicorn/run | ##### 2025-06-03T17:16:57,830819854+00:00 | gunicorn/run | 2025-06-03T17:16:58,109053012+00:00 | gunicorn/run | A required package azureml-inference-server-http is missing. Please install azureml-inference-server-http before trying again /bin/bash: /azureml-envs/azureml_c56922d4746bdea5c7cf70c56766041f/lib/libtinfo.so.6: no version information available (required by /bin/bash) 2025-06-03T17:16:58,113318961+00:00 - gunicorn/finish 100 0 2025-06-03T17:16:58,114836385+00:00 - Exit code 100 is not normal. Killing image. ChatGPT said: We have the answer: your deployment is crashing because of this line in the logs: >A required package azureml-inference-server-http is missing. -- ===== What Happened ===== Your env.yml defined all the ML packages you needed (pandas, scikit-learn, etc.), but it didn't include the Azure ML inference server runtime, which is required to serve your score.py via Azure ML Managed Online Endpoints. -- ===== Fix ===== Replace Your env.yml With This: ===== <code → ymlCopyEditname: collectibles-env channels: - defaults - conda-forge dependencies: - python=3.9 - pip - pip: - azureml-inference-server-http - pandas - scikit-learn - joblib </code> > This includes azureml-inference-server-http, the web server that Azure ML needs to route input to your run() function. -- ===== Next Steps ===== - Replace your env.yml with the version above - Rebuild and redeploy with your current deployment script - Watch the logs — this time it should say:<code → luaCopyEditModel loaded successfully. </code> -- Would you like me to: * Repackage everything with the fixed env.yml * Regenerate your deploy_to_azureml.py to be extra

safe Just say the word and I'll generate a clean ZIP or script update. == You said: ==
.....Traceback (most recent call last):
File "/Users/don.dehamer/Azure/./deploy_to_azureml.py", line 70, in <module>
ml_client.online_endpoints.begin_update(AttributeError: 'OnlineEndpointOperations'
object has no attribute 'begin_update' ChatGPT said: Great news – the model deployment
likely succeeded this time, and now you're just hitting a small SDK version mismatch. --
==== [Error: ===== <code → textCopyEditAttributeError: 'OnlineEndpointOperations'
object has no attribute 'begin_update' </code> ===== [Root Cause ===== You're using
Azure ML SDK v2, but begin_update() was used – that's from the older v1-style SDK or
incorrect documentation. -- ===== [Fix ===== To set the default deployment in SDK
v2, use this: ===== [Replace: ===== <code →
pythonCopyEditml_client.online_endpoints.begin_update(endpoint_name=endpoint_name,
default_deployment_name="blue").result() </code> ===== [With: ===== <code →
pythonCopyEditml_client.online_endpoints.update(name=endpoint_name,
defaults={"deployment_name": "blue"}) </code> -- ===== [Updated SDK v2 Flow
Summary ===== Here's the correct v2 pattern: <code → pythonCopyEdit# Create
endpoint ml_client.begin_create_or_update(endpoint).result() # Create deployment
ml_client.begin_create_or_update(deployment).result() # Set default deployment
ml_client.online_endpoints.update(name=endpoint_name,
defaults={"deployment_name": "blue"}) </code> -- Would you like a final fully
corrected deploy_to_azureml.py or bundled ZIP with everything tested and fixed? ==
You said: == defaults={"deployment_na yes py file ChatGPT said: [Your fully corrected
Azure ML deployment script is ready: [Download deploy_to_azureml.py ===== [Key
Fixes Included: ===== * Uses Standard_DS3_v2 for stability * Includes azureml-
inference-server-http via your conda env * Uses a globally unique endpoint name *
Correctly sets the default deployment using SDK v2's .update(...) method Let me know if
you'd like the final bundle as a ZIP or help testing the endpoint with a request.