Charting Your AI Journey

A Roadmap for Supervised, Unsupervised, and Generative Learning through Machine Learning and Deep Learning

Ryan Paul Lafler

Miguel Angel Bravo

Premier Analytics Consulting, LLC

Premier Analytics Consulting, LLC

www.Premier-Analytics.com

Objectives and Main Concepts

Navigating the AI Landscape: Key Concepts and Foundations

Supervised Learning: The Engine Behind Predictive Workflows

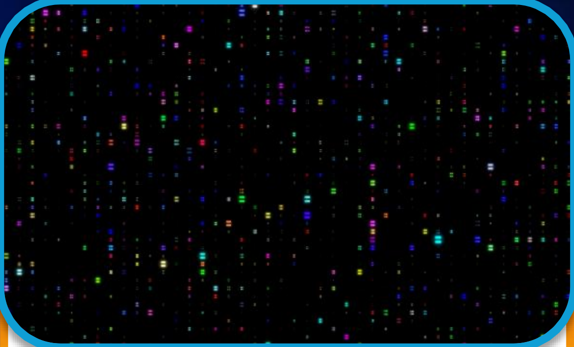

Driving Data Discovery with Unsupervised Learning

Exploring the Next Frontier with Generative AI Workflows

Premier Analytics Consulting, LLC

Navigating the AI Landscape: Key Concepts and Foundations

Defining Artificial Intelligence (AI), Machine Learning (ML), and Deep Learning (DL) and how these domains shape analytics in 2025.

AI vs ML: Understanding the Difference

- **Artificial Intelligence (AI)** is about teaching machines to mimic human decision-making
  - Examples: rule-based chatbots, game NPCs, or search algorithms
- **Machine Learning (ML):** A subset of AI where systems learn patterns from data
  - Focus areas: Predictive analytics, data discovery, and robust training & learning
- Domains of Machine Learning:
  - **Supervised Learning**: Uses labeled data to predict outcomes
  - **Unsupervised Learning**: Finds patterns in unlabeled data
  - **Semi-Supervised Learning**: Combines small labeled sets with large unlabeled data for training

Advantages of AI Systems Built from ML

- **Dynamic Learning**: ML-powered AI adapts rules and models automatically as new data becomes available, reducing manual re-programming
  - Codebases do *not* need to be manually re-defined → automatically done
- **Pattern Discovery**: Detects hidden trends, anomalies, and insights that humans might miss, even in complex datasets
- **Scalability**: Easily scales to large, constantly evolving data sources without needing static decision rules
- **Continuous Improvement**: Each interaction or new dataset attempts to improve and learn from previous mistakes

Scope-Defined AI vs. Generative AI

Traditional AI

- **Task-Specific Models:** Built for narrow tasks (classification, recommendation)
- **Fixed Historical Data**: Retraining needed for updates; models rely heavily on curated datasets
- **Manual Feature Engineering:** Requires experts to select features and rules
- **Domain-Specific Expertise:** Requires deep subject matter knowledge to tune algorithms

Generative AI

- **Multi-Tasking & Reusability:** Adapts to many tasks with minimal fine-tuning
- **Embeddings for Memory**: Stores context for retrieval, augmentation, and reasoning
- **Creative Freedom (Temperature)**: Generates unique, dynamic responses (synthetic data)
- **Zero-Shot Training**: Model can perform tasks not trained on by leveraging existing info
- **Reasoning and RAG**: Context & justifications

How Deep Learning (DL) Powers Generative AI

Deep learning (DL) powers the leap to Generative AI by automatically learning representations, scaling for big data, and enabling transfer learning, multimodal reasoning, and creative synthesis through neural networks and transformers.

- **Representation Learning:** Neural networks automatically learn hierarchical features from raw data, automating feature engineering
- **Scalability:** Trains on massive datasets with GPUs/TPUs → improves generalization
- **Transfer Learning:** Fine-tune pre-trained models for new domains / tasks with minimal data
- **Multimodality:** Processes text, images, video, & audio together for cross-domain reasoning
- **Transformer Architectures:** Attention mechanisms for long-range context & reasoning

Multimodal Generative AI in Action

- **Multimedia and Entertainment:**
  - Generate scripts, music, and visuals; upscale images; automate captions and dubbing
- **Summarization and Translation**:
  - Summarize text, translate languages, and write rich image descriptions
- **Retail and E-Commerce**:
  - Deliver personalized recommendations, intelligent shopping assistants, and agents that summarize and answer questions personalized to any product
- **Software and IT**:
  - Generate, refactor, and debug code; automate documentation and database queries

Multimodal Generative AI in Action

- **Life Sciences and Healthcare:**
  - Analyze scans, generate synthetic medical data, and explain patient outcomes
- **Meteorology:**
  - Downscale forecasts, deliver localized real-time weather summaries, and power severe weather alert systems
- **Finance and Risk Management**:
  - Fraud detection, risk assessment, automated compliance reporting, financial forecasting, personalized insurance policy recommendations, and synthetic data generation for model training

Key AI Domains Used in Industry and Research

Supervised ML

- Pre-programmed algorithms for **classification**, **regression**, **forecasting**, **ranking**, and **survival analysis** (*with censoring*)
- Requires **labeled data** with set of predictors *mapped* to target feature(s) given a **loss function** (learns from ground truth)
- Algorithms possess **hyperparameters** that define *structure* of supervised algorithms → must be optimized

Unsupervised ML

- Pre-programmed algorithms for **data compression**, **reduction**, **clustering**, **similarity scoring**, and **association analysis**
- Can mine **raw** and **unprocessed data** to reveal structure, extract insights, and uncover patterns → no ground truth necessary
- Algorithms contain **hyperparameters** that are chosen based on domain knowledge or some metric

Generative AI

- Builds on ML and DL to **generate new data**, **content**, or **patterns** (e.g., text, images, code, simulations)
- Excels at **retrieval-augmented generation (RAG)**, **contextual understanding**, and **creative synthesis**
- Adapts to multiple domains with **prompt engineering, transfer** learning, and **layer fine-tuning** → reduces re-training needs

Premier Analytics Consulting, LLC

Supervised Learning: The Engine Behind Predictive Workflows

Leveraging labeled datasets, hyperparameter optimization, and diverse algorithms for classification, regression, and ranking in production.

Supervised ML and Predictive Analytics

- Supervised algorithms train from labeled data, **iteratively improve** and adjust **internal parameters (weights)** to minimize some objective (loss) function

- **Assumptions** and **hyperparameters** constrain the flexibility of ML algorithms and define their internal structure
  - **Linear Regression** assumes dependent (target) variable is normally distributed
  - **Decision Tree** makes *no* assumptions about data (very flexible algorithms)
    - Model non-linear relationships and can be constrained by **hyperparameters**
  - **Support Vector Machine (SVM)** can model *high dimensional* data with few observations (wide datasets)

Mitigating Overfitting & Underfitting

- Supervised algorithms require *labeled* data to learn from the **ground truth**
    - Predictors mapped to target feature(s) by splitting data into *training* and *testing* sets
    - Algorithms are trained on training set; testing set used for *evaluation*
- **Overfitting**: Model is too sensitive (**too flexible**) to changes in data and *quickly adapts*
    - Referred to as **high variance** → Hyperparameters fail to regularize model complexity
- **Underfitting**: Model is too rigid (**too constrained**) and does not adapt to changes in data
    - Referred to as **high bias** → Assumptions / hyperparameters are too restrictive

Tradeoff Between Interpretability & Complexity

- **Model complexity increases flexibility and accuracy but reduces interpretability, making it harder to explain decisions**
  - **Multiple Linear Regression** allows us to understand effects of predictors on target feature
    - Can evaluate partial slopes (weights) of each factor, interaction term, or higher-order term to know effects on target feature → we assume linear relationship
  - **Decision Tree** is branching if-else structure → tree splits on best feature to make prediction
    - No assumptions on data → can model non-linear relationships while still being interpretable (middle-ground)
  - **Neural Networks** can model complex relationships but are “black boxes”

Charting Your Supervised ML Roadmap

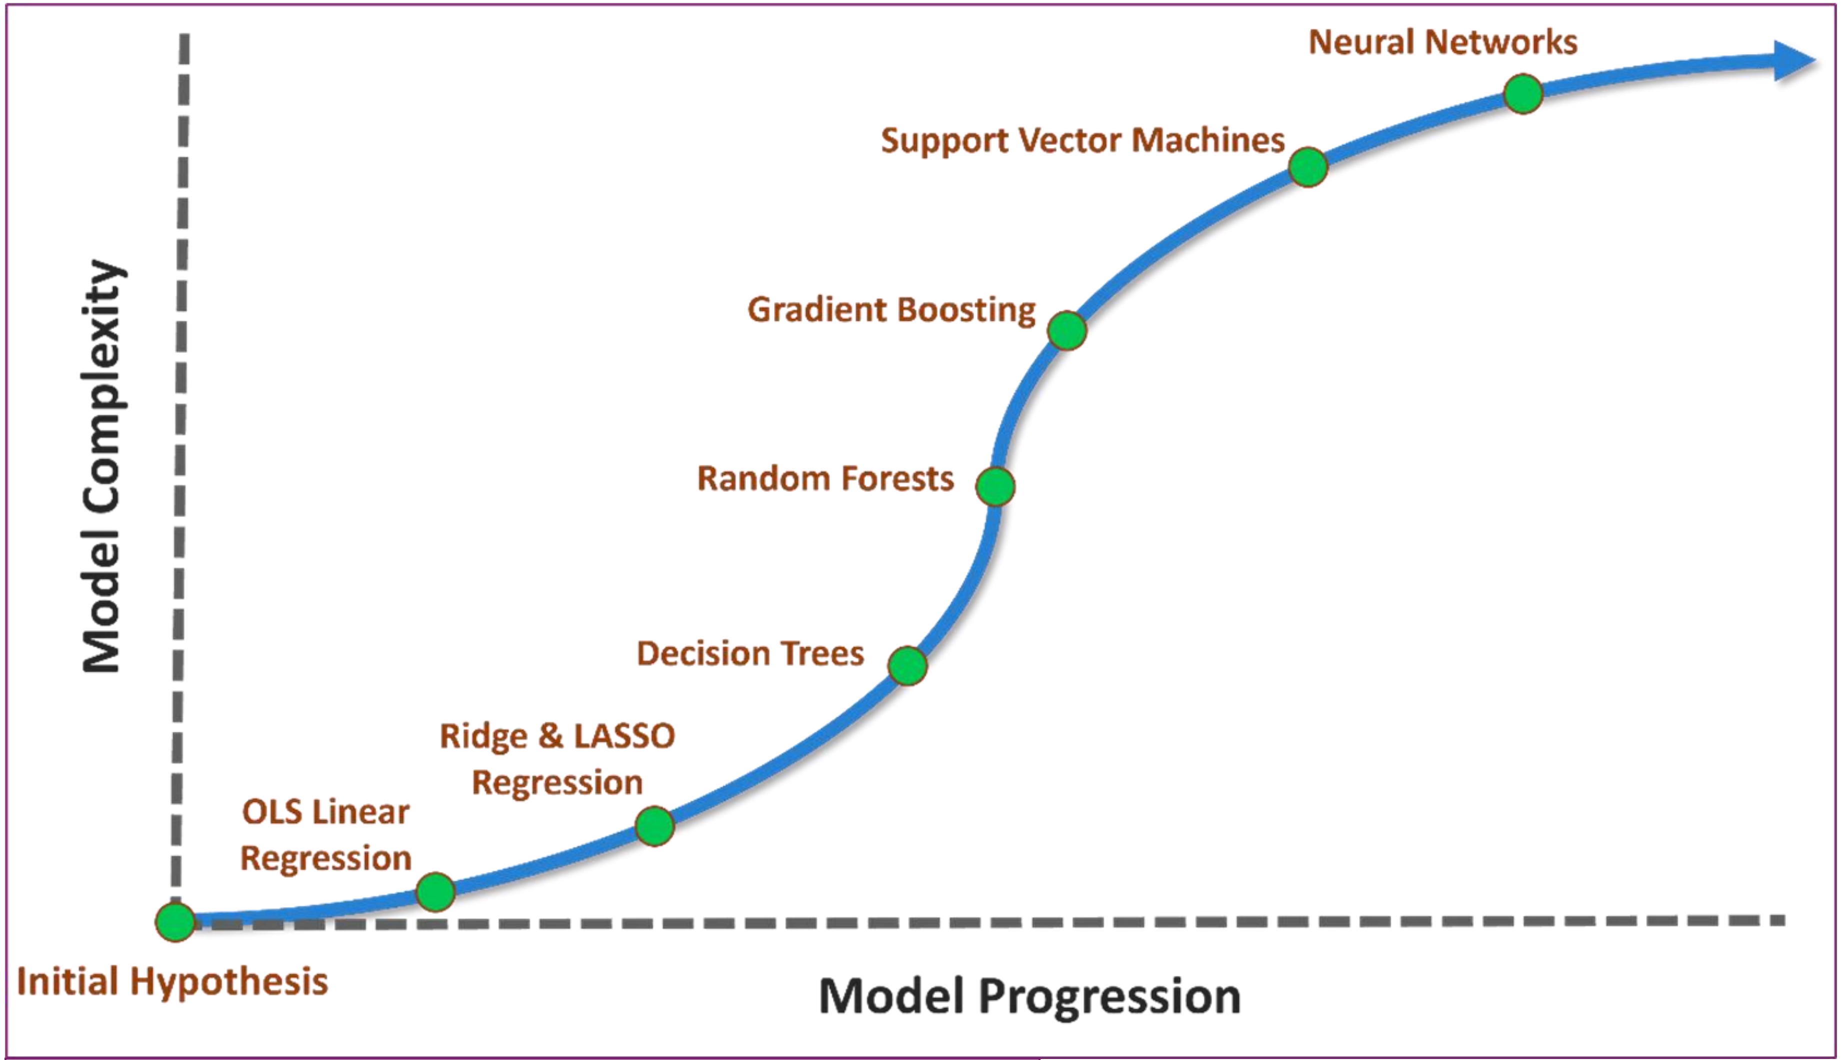

Ryan Paul Lafler | *Premier Analytics Consulting*

Hyperparameters of a Decision Tree

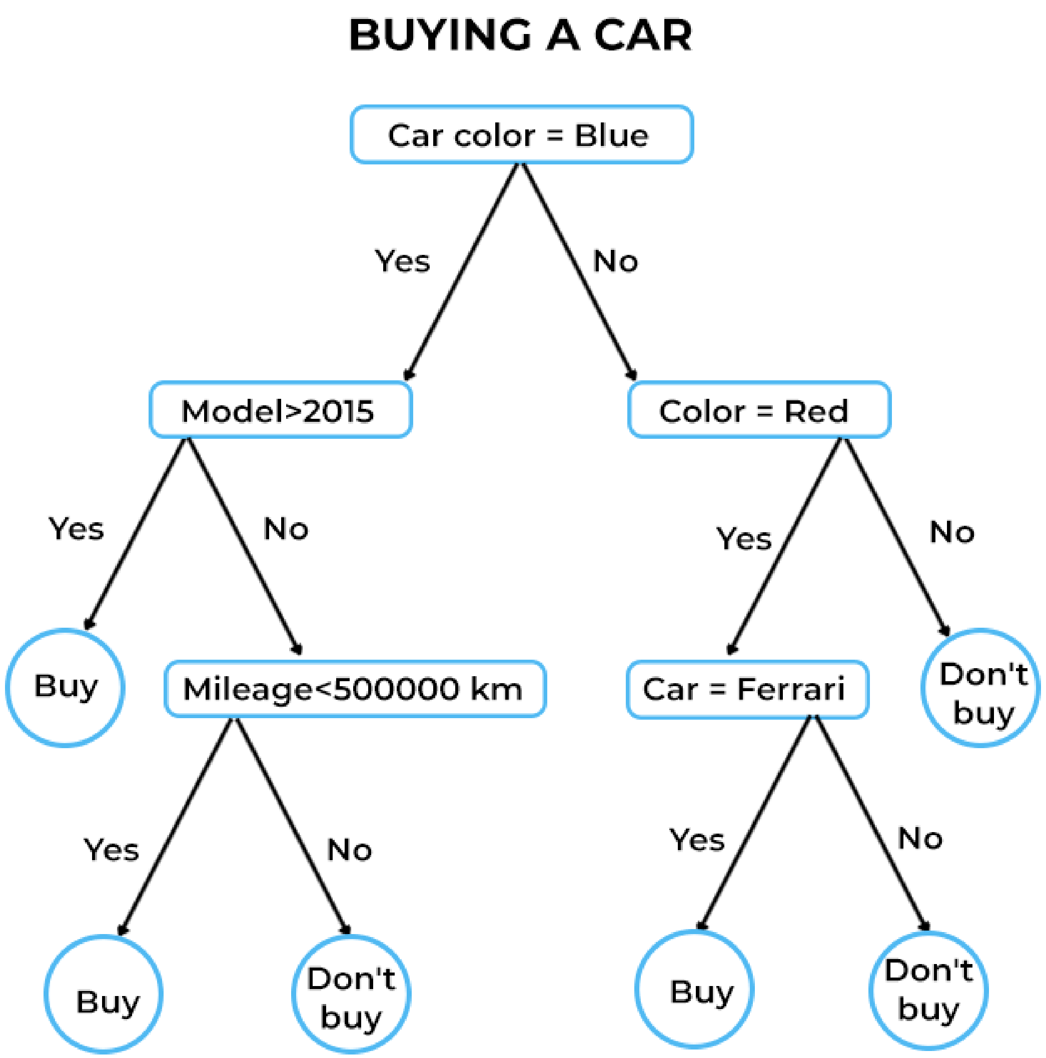

Premier Analytics Consulting, LLC

Driving Data Discovery with Unsupervised Learning

Training algorithms to group, compress, detect anomalies, reveal patterns, and uncover relationships on unprocessed data without labels.

Data Discovery with Unsupervised ML

- **No Labels Required:** Learns patterns directly from raw, unprocessed data → no target labels or explicit outputs needed
- **Insight Generator:** Finds hidden structures, trends, and relationships in large, unlabeled datasets → many techniques adapted from **signal processing**
- **Signal from Noise**: Isolate and identify signal (attributes giving most variance) from randomness (noise) inherent to all data structures
- **Data-Hungry Approach:** Requires substantial volumes of raw data to uncover reliable and meaningful insights
- **Representation Extraction**: Encode features into more processable inputs to models

Clustering: K-Means, Hierarchical, & DBSCAN

- **Clustering:** Group data points by similarity to uncover structure in unlabeled datasets
  - **K-Means**
    - Partitions data into k clusters by minimizing distance to cluster centroids
    - Simple, fast, and works well with spherical clusters
  - **Hierarchical**
    - Builds a tree of clusters (dendrogram) by iteratively merging or splitting groups
    - Great for visualizing data structure, relationships, & nested clusters
  - **DBSCAN**
    - Finds clusters based on data density; identifies outliers as noise
    - Works well for irregularly shaped clusters and uneven densities

Hierarchical Clustering using a Dendrogram

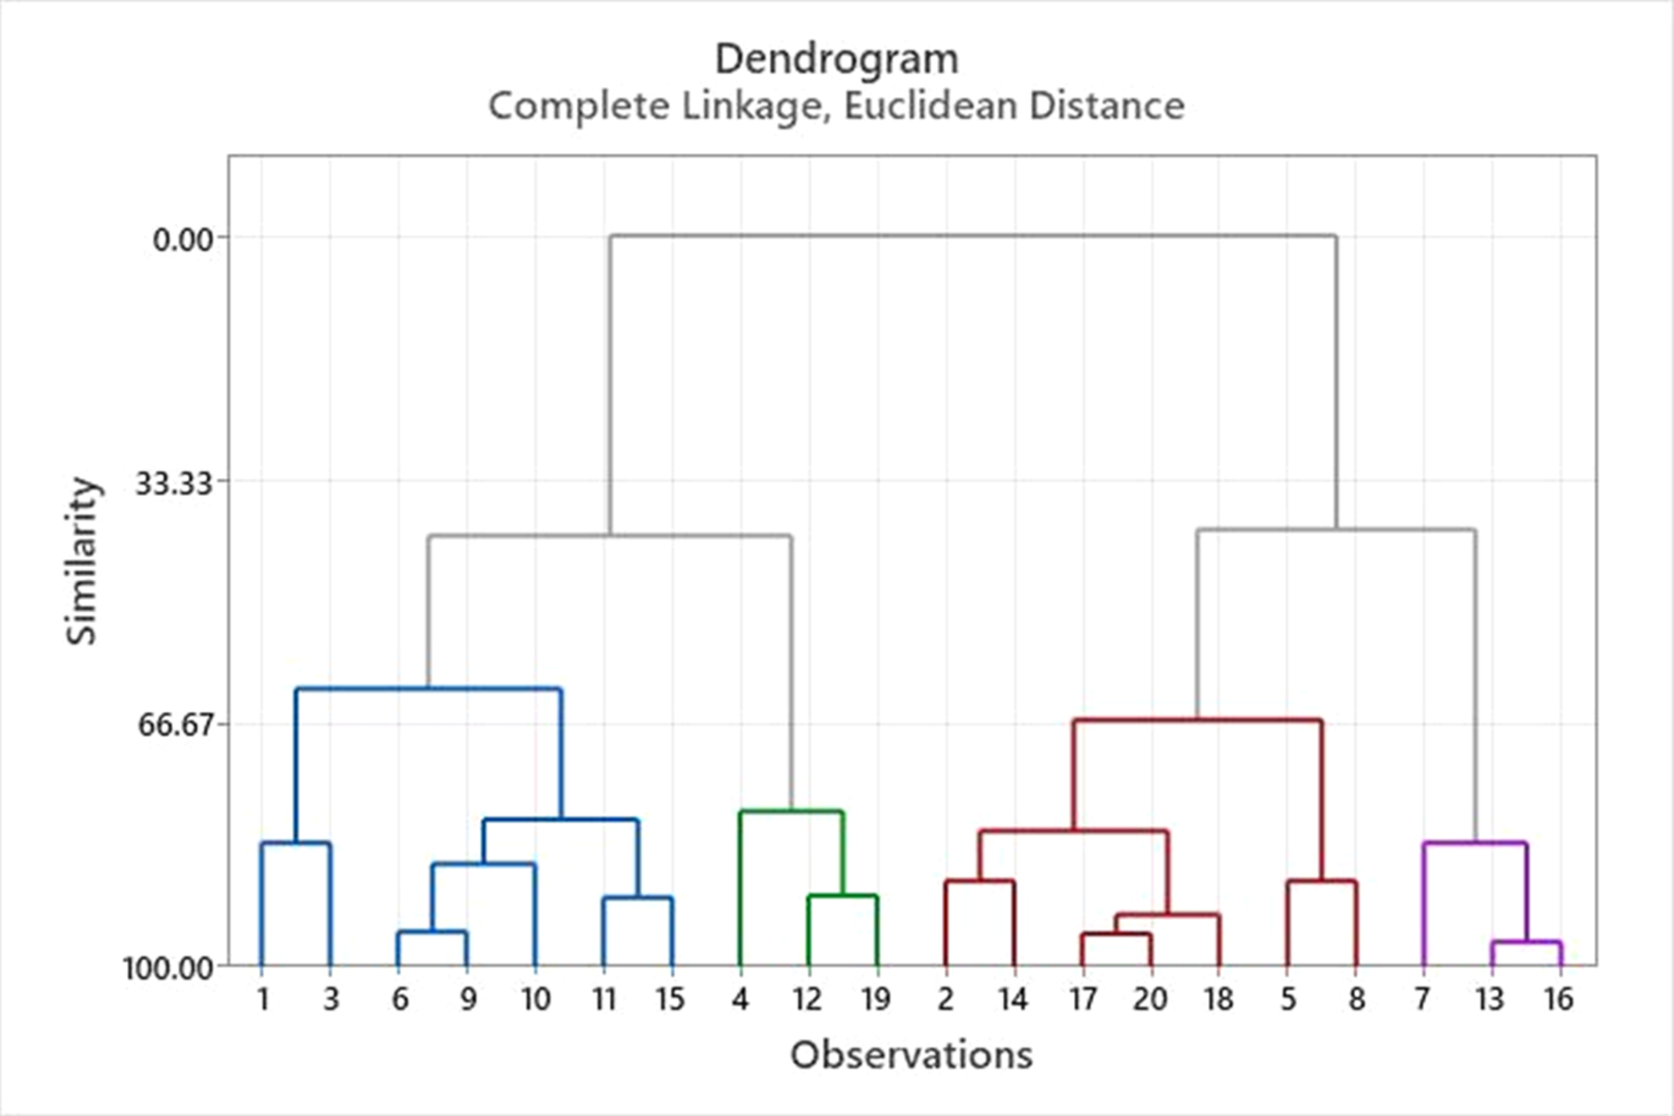

K-Means Clustering with Centroids

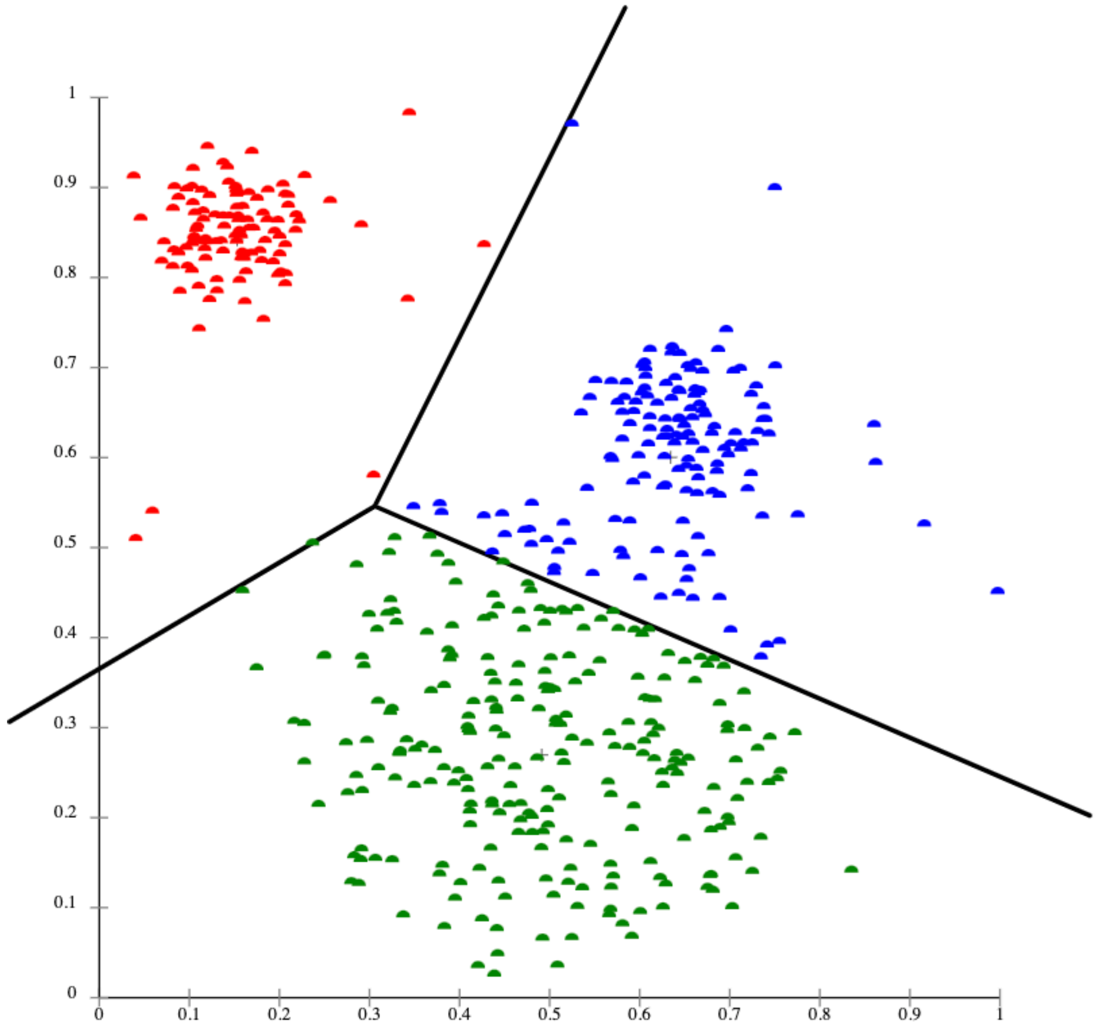

DBSCAN Clustering to Uncover Distributions

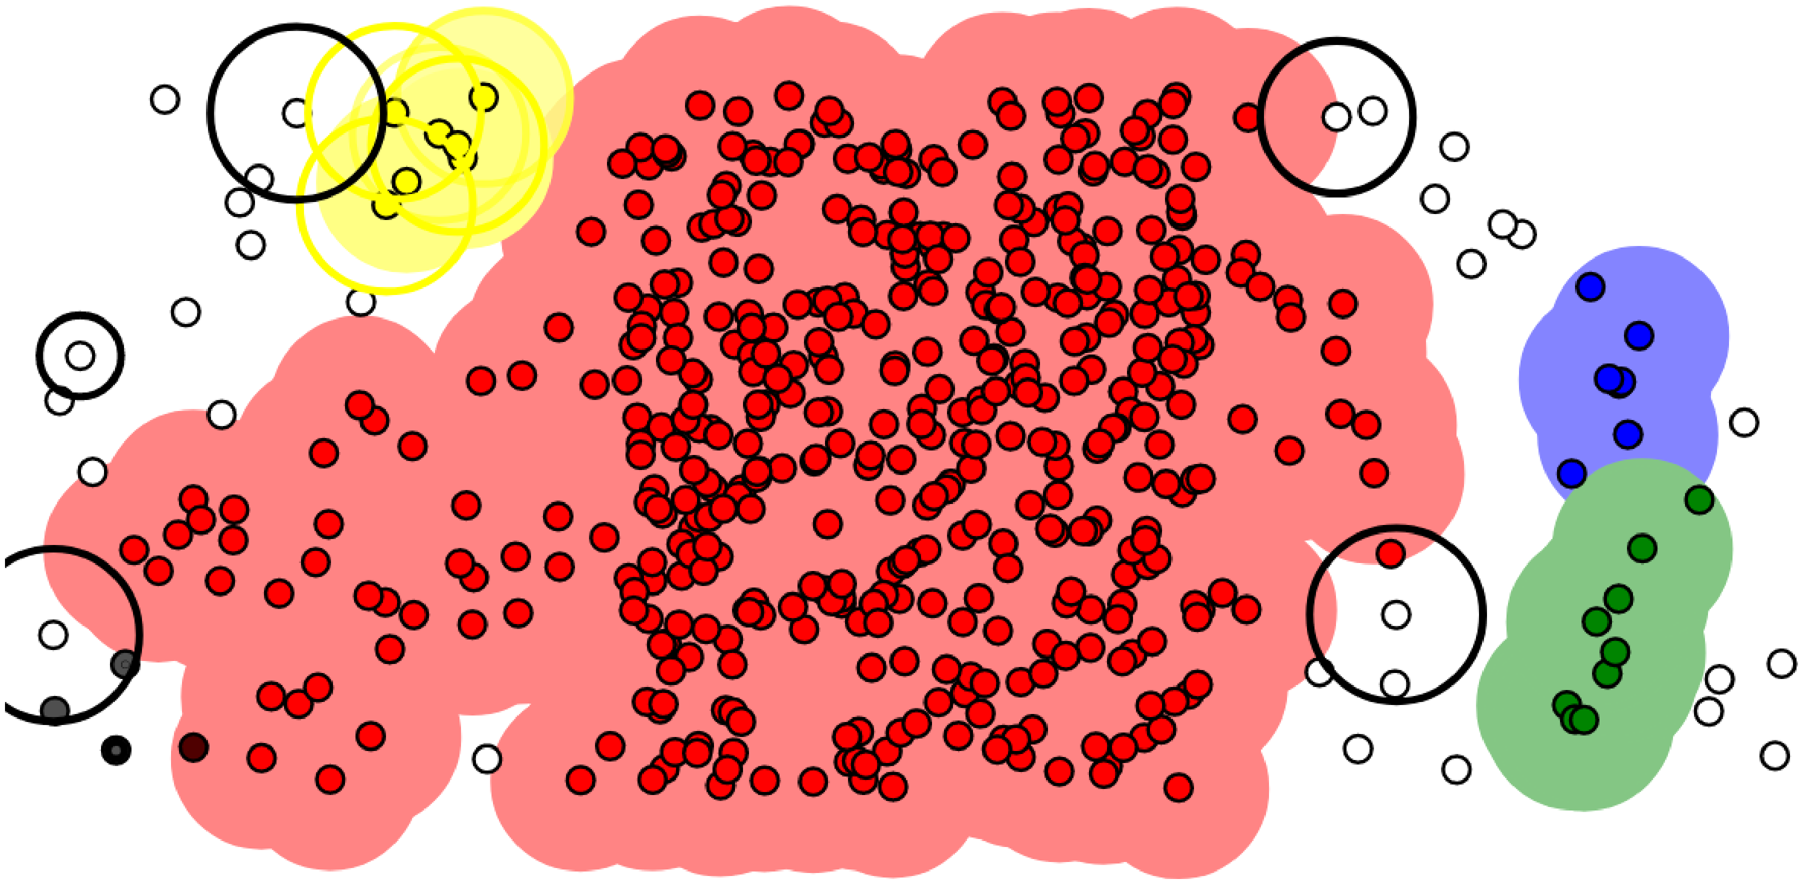

Dimensionality Reduction & Distribution Modeling

Dimensionality Reduction

- Reduces high-dimensional data into key components while preserving structure
- **PCA:** Captures maximum variance directions for visualization & compression
- **LLE & IsoMap:** Preserve non-linear manifold structure for better pattern discovery
- Enables **clustering** & **pattern recognition** by **removing noise and redundancy**

Gaussian Mixture Models (GMMs)

- Models datasets as a mixture of multiple Gaussian distributions (*number is unknown*)
- Provides **probabilistic cluster assignments** (*soft clustering*) for greater flexibility
- Unlike K-Means, assigns probabilities (not distance metric) for membership in multiple clusters
- Learns underlying **density of data** → enables **simulation** and **synthetic data generation**

Dimensionality Reduction & Data Compression

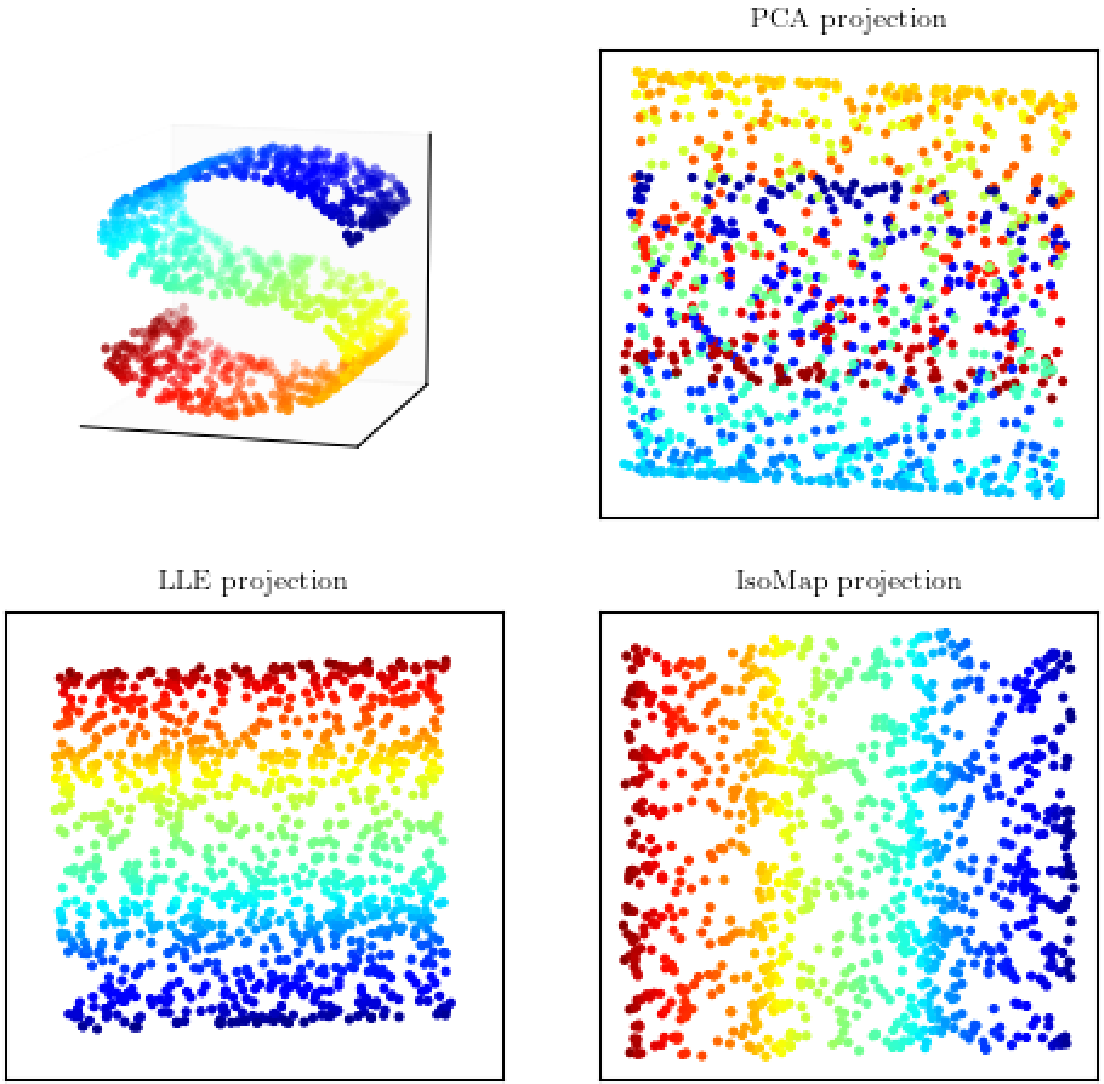

Gaussian Mixture Model (GMM) in 3D-Space

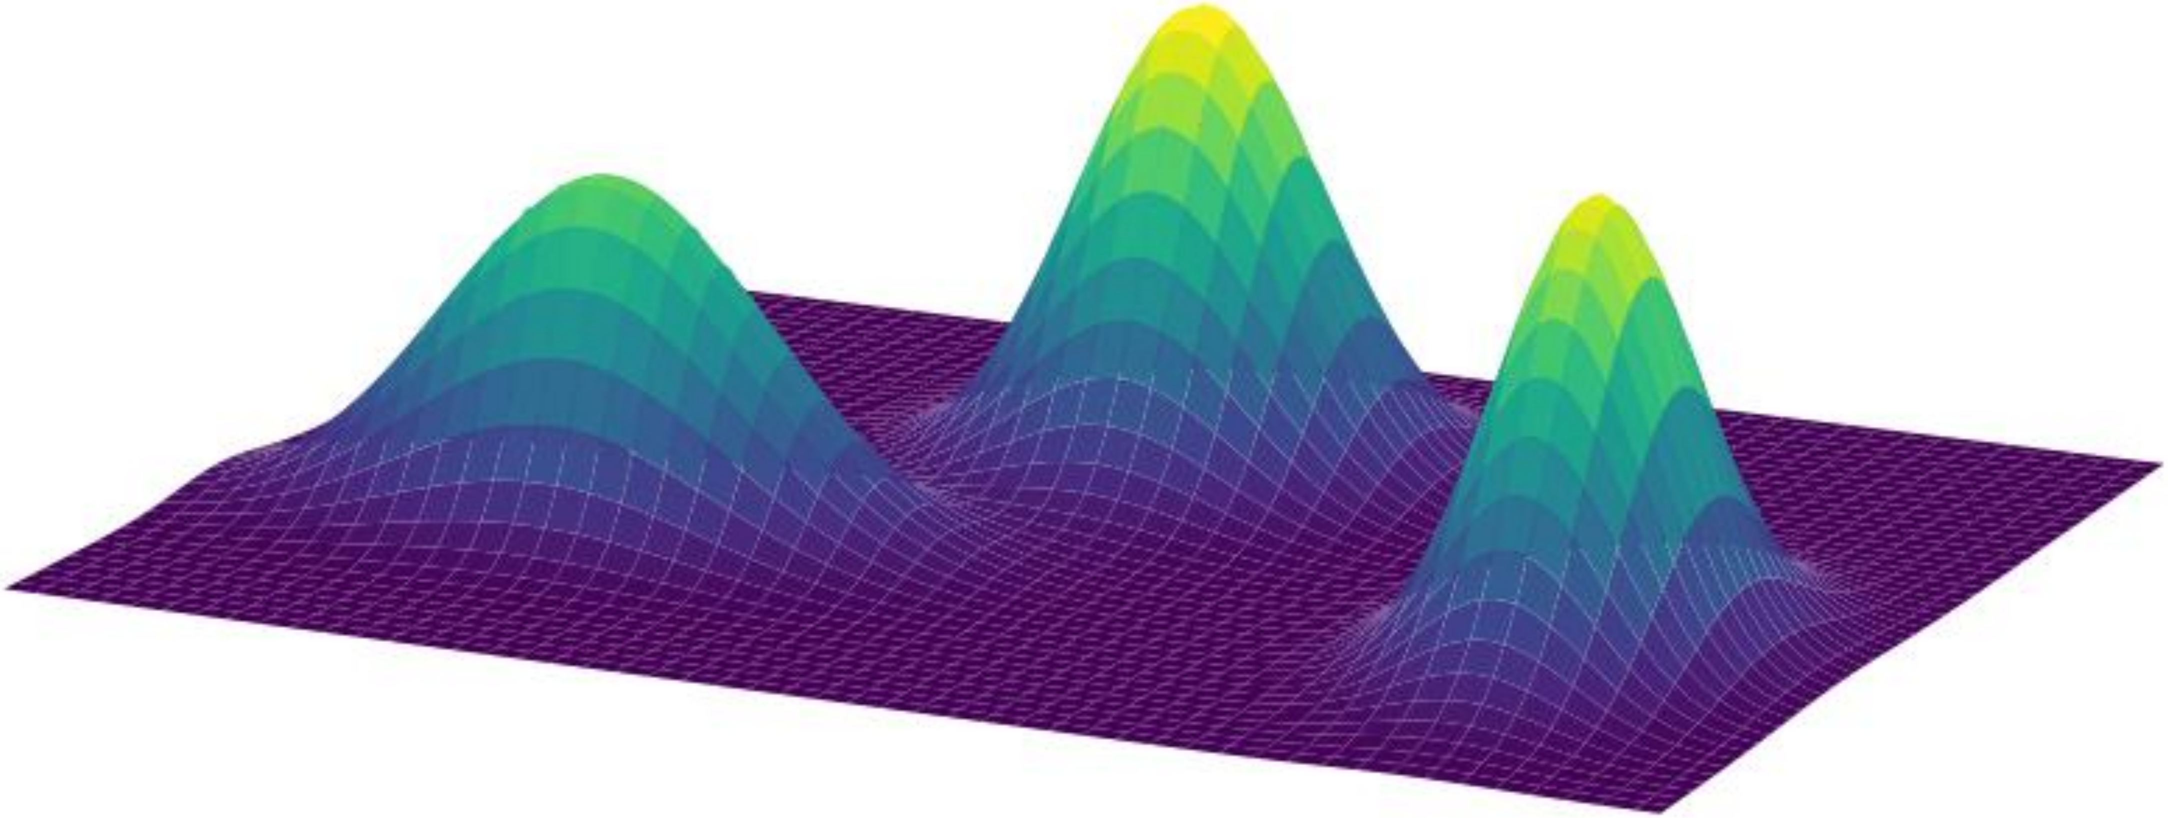

Premier Analytics Consulting, LLC

Exploring the Next Frontier with Generative AI Workflows

Harnessing encoder, decoder, and hybrid transformer models to create content, capture context, and enable reasoning and information retrieval.

Generative AI Architectures

- **Transformers** key to generative AI workflows → parallelize **attention** layers in neural networks and replace sequential networks such as RNNs and LSTMs
- **Encoder-Only architectures** excel at representation learning, embeddings, and retrieval; power search, classification, and semantic similarity
  - Create contextual representations (vector embeddings) from input tokens
  - Examples include **BERT** and **RoBERTa**
- **Decoder-Only architectures** provide autoregressive generation for text, code, simulations, and creative tasks; optimized for long-form output
  - Each output token *only* focuses on previous token(s), does not see ahead of current token → examples include GPT family & LLaMa family

Generative AI Architectures

- **Encoder-Decoder architectures** are ideal for translation, summarization, multimodal reasoning; integrate comprehension + generation
  - Encoder-part transforms input sequence to vector embeddings
  - Decoder-part consumes vector embeddings and generates output sequence
    - Encoder-part builds a rich representation by reading entire input sequence and transforming into latent (vector) space
    - Decoder generates tokens **one-at-a-time** using **casual attention** to maintain sequence order
  - Unlike Decoder-Only architectures, Encoder-Decoder transformers excel at summarization and translation tasks by focusing generation guided by input context

Retrieval-Augmented Generation (RAG)

- **Bridges search + generation:** Combines external knowledge retrieval with large language models (LLMs) to improve factual accuracy
- **Retriever module:** Fetches relevant documents or embeddings from a vector database or search index based on the query
- **Vector Databases:** Purpose-built for storing and searching vector embeddings
  - Enables fast similarity search and semantic retrieval to power RAG, recommendations, and multimodal AI
- **Generator module:** Conditions its response on both the input query and retrieved content, ensuring grounded outputs
- Good for **Q/A systems**, **chatbots**, & **technical documentation search**

Retrieval-Augmented Generation (RAG)

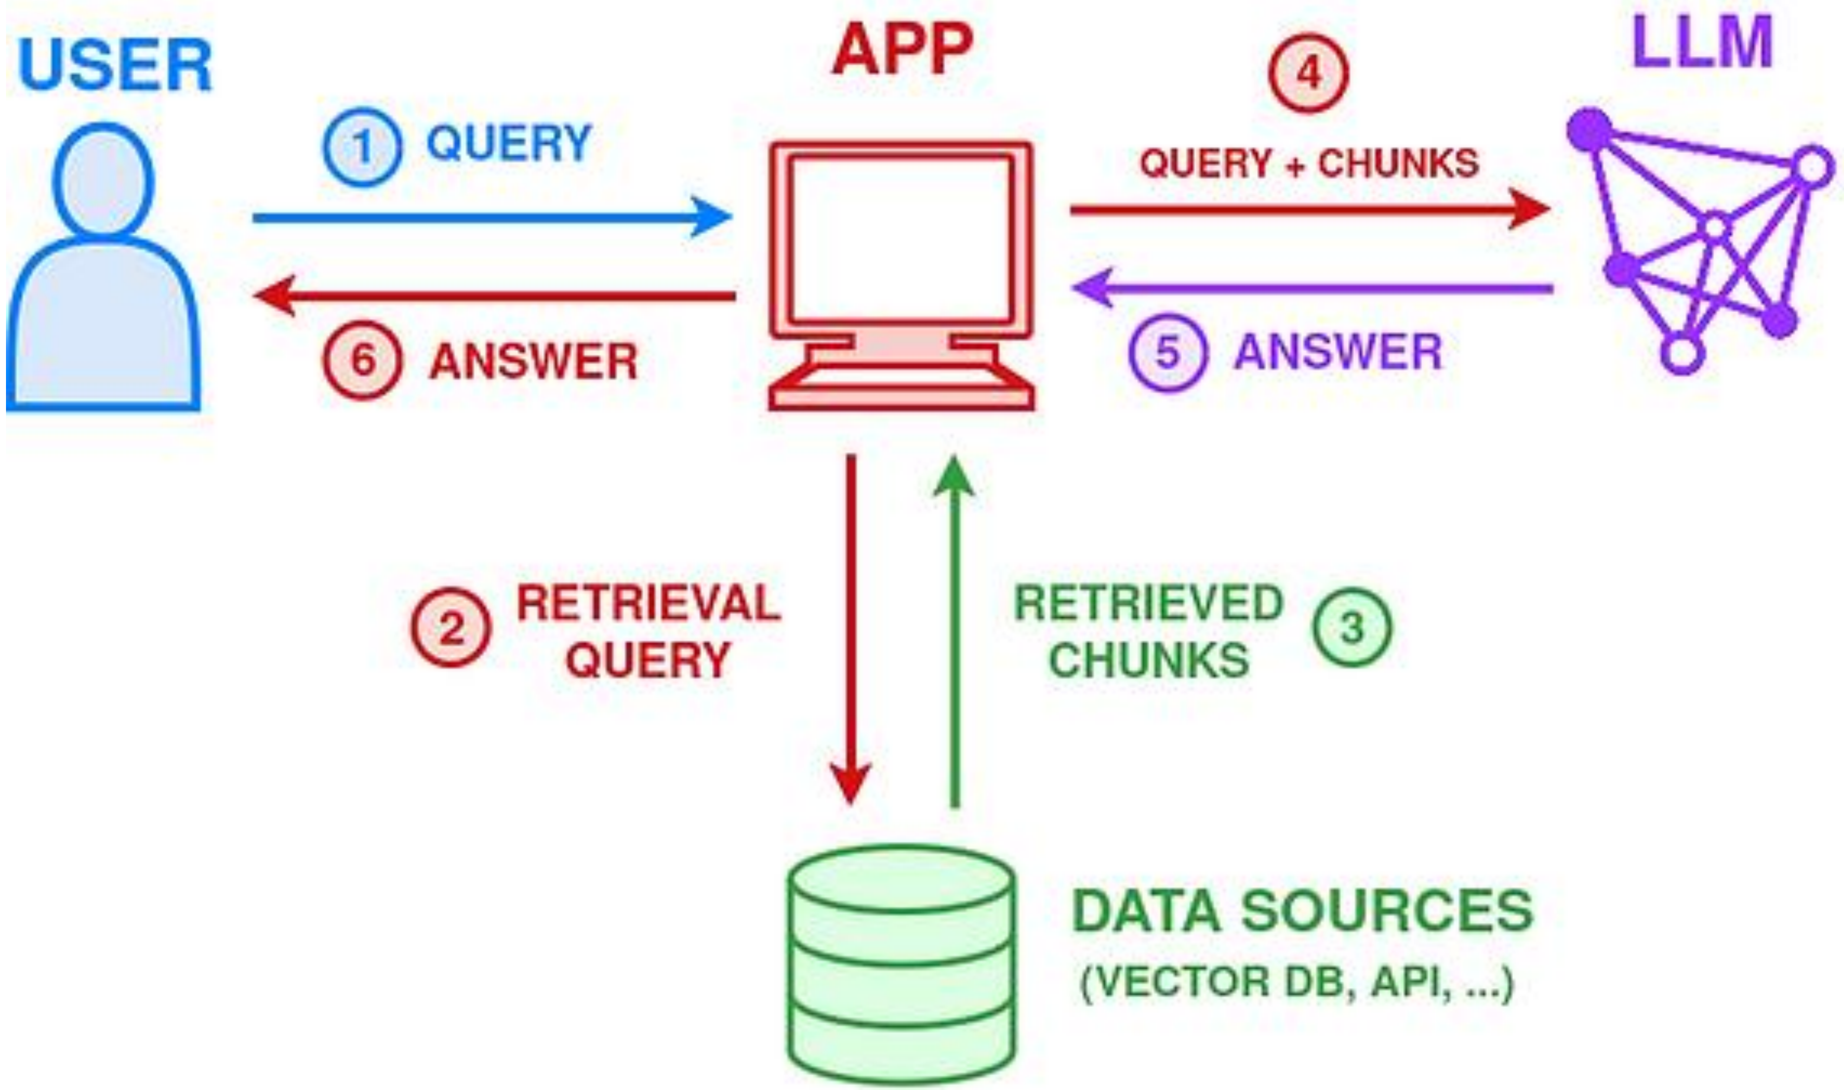

Agentic AI Workflow

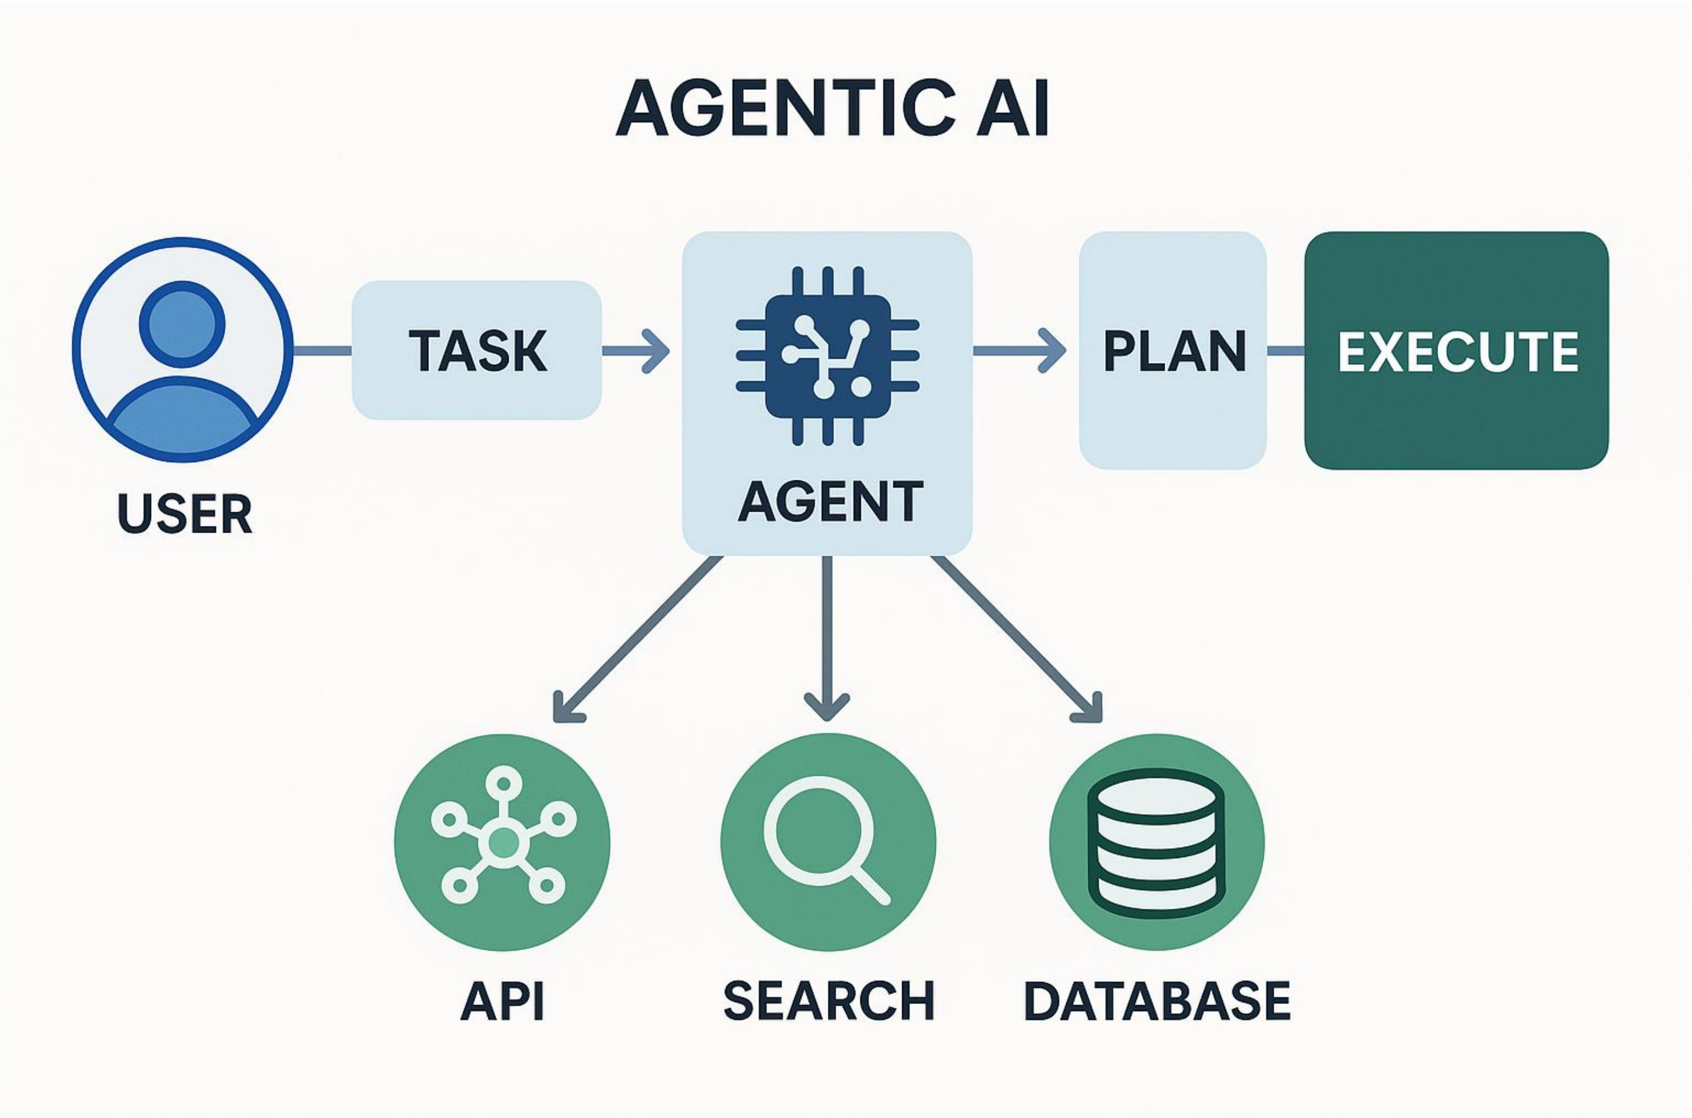

Charting Your AI Journey

A Roadmap for Supervised, Unsupervised, and Generative Learning through Machine Learning and Deep Learning

Ryan Paul Lafler
rplafler@premier-analytics.com

Miguel Angel Bravo
mabravo@premier-analytics.com

Premier Analytics Consulting, LLC

www.Premier-Analytics.com