AI: A Leader's Guide

Prompt (from another leader): *"What are you doing with AI?"*

You typing into ChatGPT: *"What am I doing with AI?"*

The exchange above is tongue in cheek, but a version of it is happening everyday. This short paper is geared towards arming you with what we are seeing done in **real-world AI implementations** today.

The AI conversation is evolving quickly. It is hard to keep up. It helps to start with some basic building blocks. The **2 big topics to build on are Generative AI and Machine Learning**. Most of the news cycle today is dominated by advances in Generative AI, GenAI for short. In today's enterprises, there are great use cases for both.

Generative AI

What is it? Algorithms that can create new content, such as text, images, and music. Different algorithms/models are better for specific tasks. For example, code generation vs generating a marketing image

Leaders:
- OpenAI ChatGPT / Microsoft Copilot
- Google Gemini
- xAI Grok

Use Cases:
- Reading and summarizing reviews
- Creating graphics for a product catalog
- Answering questions based on a knowledge base

Common Challenges:
- Grounding and hallucination
- Leaking proprietary information to the public
- LLM selection, fit for purpose
- Agentic framework competition

Machine Learning

What is it? Algorithms that can learn from and make predictions based on data. ML is instrumental in analyzing large datasets, identifying patterns and making predictions about what might happen

Leaders:
- Microsoft (AI Investment and Research)
- Google (TensorFlow)
- Amazon (SageMaker)

Use Cases:
- Predict when a piece of equipment will be available
- Estimate the size of an oil reserve
- Cluster customers based on behaviors

Common Challenges:
- Trust gap (model performance vs human intuition)
- Data quality, quantity and labeling
- Model selection and optimization
- Access to qualified talent

What are the common challenges we see in the marketplace today:

Science experiments. Many consultants are learning on their customer's dollar working on poorly conceived and executed efforts that have dubious real-world value

Not knowing where or how to start is often one of the biggest road-blocks for customers. Starting with a solid foundation is how to make continuous, material improvements in an AI journey

Many customers **do not understand the core value** and premise of AI. A well-formed question needs to be created before looking for an answer

Operational trust. Making predictions and building models is not as difficult as it used to be. However, making the leap from a working model to production application takes a partnership and domain expertise

Serverless Solutions

Email: Contact@serverless-solutions.com

Phone: (281) 810-6270

AI: A Leader's Guide

The gravity of the conversation today is around GenAI so let's discuss what is emerging now and what you should pay attention to.

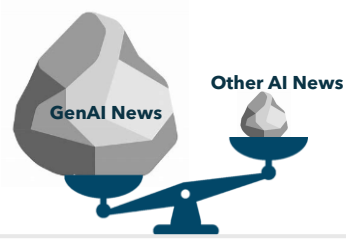

Putting the pieces together

Getting the easy win by using something like Gemini or Copilot in your productivity suite is no longer what will keep you at parity with your peers in the market. You should be thinking about parts of your processes that can be encapsulated as an agent, then how can I make that agent more valuable by extending it's reach with things like MCP then finally how does that agent interact with a broader process using agentic frameworks.

There is a lot to consider, but the opportunity to impact your business in a powerful way is undeniable.

Multi-Modal Models

These models can understand text, images, video and sound. This opens many new use cases and makes interacting with the models more intuitive.

Agents

Autonomous tool capable of making independent decisions and executing tasks with out human intervention.

Agentic Frameworks

Now that you have an agent, how do you get your agent to talk to other agents and participate in a more complex process? That is where agentic frameworks come in.

Model Grounding

Generative AI is guessing what the user wants to hear and can be creative with responses, often responding to the same question differently. What if I want the responses to be purely factual and consistent? That is where model grounding and RAG style solutions come into play.

Custom or Package Solutions

Software vendors have raced to include GenAI into their products. Knowing when to use a product vs build you own solution requires experience in both. Just try to navigate the term "Co-Pilot". Is it a website? A part of Office365? A service in Azure? The answer is yes. How do you know which one you need?

MCP

The Model Context Protocol (MCP) goal is to enable seamless interaction with external tools, data sources and APIs.

The hardest part of this conversation is the imagination it takes to think of ways this can impact your business. However, closely behind that are the technical considerations for how this ecosystem works. For example, if I have an agent that has access to healthcare information, what information is the agent allowed to share and with what other agents or other consumers? These considerations get complex quickly when many agents are involved and when agents start crossing boundaries like communicating with an agent or service that is not part of my IT landscape.

Generative AI Key Moves:

1. Secure your IP:
When you put information into public Generative AI platforms, it becomes that company's property and can be shared in future responses to anyone. You can protect yourself with the right platform and configurations.

2. Turn potential into profit:
When a real-world use case is identified, simplify it until it works well and can generate real world cost improvements or top line revenue growth. Then go deep with it. Operationalize it, grow it and put those learnings into new use cases.

3. Invest in an eco-system:
There are many agents, frameworks, models and supporting infrastructure. Pick an eco-system and build your organizational muscle with it. Investing in your team but partnering with experts that have delivered production solutions already.

Serverless Solutions

Email: Contact@serverless-solutions.com
Phone: (281) 810-6270

AI: A Leader's Guide

Now let's talk about the less popular AI sibling, Machine Learning.

While Machine Learning lacks the magic Generative AI has in mimicking human intelligence, it makes up for it with raw ability. That raw ability is math and **using math with lots of data to predict the future**. Machine learning can be applied to predict tomorrow's inventory conditions, interpret an image, pinpoint when an equipment failure will happen, tell you when a customer will churn and many other future looking events.

Machine Learning efforts have been delivered, and in production operations, for years and there is a level of maturity in these efforts that is still emerging with Generative AI projects. There is also a level of certainty, science even, that comes with a machine learning project. For example, a model might predict a customer will churn but also tell you the model is 86% certain of the outcome. This precision allows implementers to work closely with domain experts and operators put these solutions into a production setting with a higher degree of confidence.

Machine Learning Key Moves:

1. Label Data:

To predict the future, you need lots of examples of something happening in the past. Many systems have this history, but it is not labeled to capture the event you are trying to predict. This is the foundation of a successful ML project.

2. Experiment:

Knowing what will create a successful model takes trial and error as well as expertise with specific types of models. Iterate quickly and invest in fundamentals like normalized data before you get carried away with algorithms. We are a fan of the many "AutoML" capabilities available today.

3. Build Trust:

Encourage dialog between operators and technologist because they both must cross the trust gap together when it comes time to put technology in a place where humans have historically operated.

Project Examples (GenAI and ML):

Operational Data + GenAI:

In this effort we combined operational data with generative AI. By using customer life-time-value with generative AI's ability to read customer reviews. Front line leaders know if a high-value customer had a bad experience and can react.

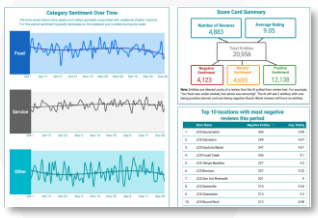

Proposal Agent:

Using an agentic framework, vector search, a modern LLM and some clever UX we created a tool for generating complex proposal summaries using a library of an organization's knowledge. This agent can work in concert with other agents to generate more complex proposal responses.

Out of Stock Alert:

In this effort we used commodity security cameras to monitor a warehouse and analyze video for low and out-of-stock conditions for a vendor managed inventory solution. Models detected when a customer would run out of stock and submitted an order to restock.

Email: Contact@serverless-solutions.com
Phone: (281) 810-6270

AI: A Leader's Guide

If you made it this far, this is where we tell you how Serverless Solutions can help you on your AI journey.

We are passionate about our customers' success. AI has been a challenge for our customers and our team because the stakes are high to make real improvements, but also because it is hard to guarantee outcomes when we are all operating on the cutting edge of technology. We have real world, proven expertise that can accelerate your AI efforts today.

Our Differentiators

Accelerators and Templates

We have delivered many successful AI solutions across Generative AI, Predictive AI and Time Series AI. All requiring different kinds of skills and expertise. Use our learnings captured in tools and templates to accelerate your AI journey.

Real-world experience

Customers do not usually appreciate the gap that shows up between a "working model" and an operations person trusting the output of a model. Bridging that gap requires an ability to simplify and build processes that allow trust to develop with new AI solutions.

Production Deployments

What makes a good AI solution at time of deployment will likely change over time. Monitoring, refining and changing your model performance is perhaps more difficult than the actual build of the model. Lean on our experience with production deployments.

Specific Domain Expertise

AI is a huge field now. Specialization is key. Specialization in a domain area like an industry, process or technology is required to be successful.

Free PoCs

If you are working on a model we are unfamiliar with, we will invest our own dollars in a proof of concept so that all parties know what is possible and what success looks like before entering into a commercial agreement to execute a project.

Education

Some of the most impactful AI conversations we have had are just simple discussions on what is possible, what kind of data you need and how AI actually works. After these conversations, customers become truly dangerous and produce use cases our team never would have thought of.

Next Steps:

1. **Kick the tires:** Sign-up for a 2-month trial of Microsoft Copilot on us
2. **Prompt-a-thon:** Join us at a Prompt-a-thon where we will do some education and have fun with Generative AI
3. **Executive Workshop:** Take a guided tour of what is possible in the context of your organization
4. **Project / PoC:** If you have a use case identified and need help, engage us to provide a proposal for your next AI effort

This can be an intimidating topic, but we love talking about it. Give us a call and we will put you in control of your AI journey.

Serverless Solutions

Email: Contact@serverless-solutions.com
Phone: (281) 810-6270