Embracing AI:

Tech to Allow you to Leave the office on Time

Municipal Finance Institute 2025

J.R. Wright

CIO @ Culver City

Bryant Baek

CEO & Co-Founder @ Bettercast AI

Erica Olsen

CEO @ Madison AI

Steve Hargis

CIO @ City of Moreno Valley

What's happening in your agency related to AI?

Exploring | Using Co-Pilot/GPT | Using Purpose-Built AI Tools

Top Ten Factors Driving AI Adoption

IBM Global AI Adoption Index

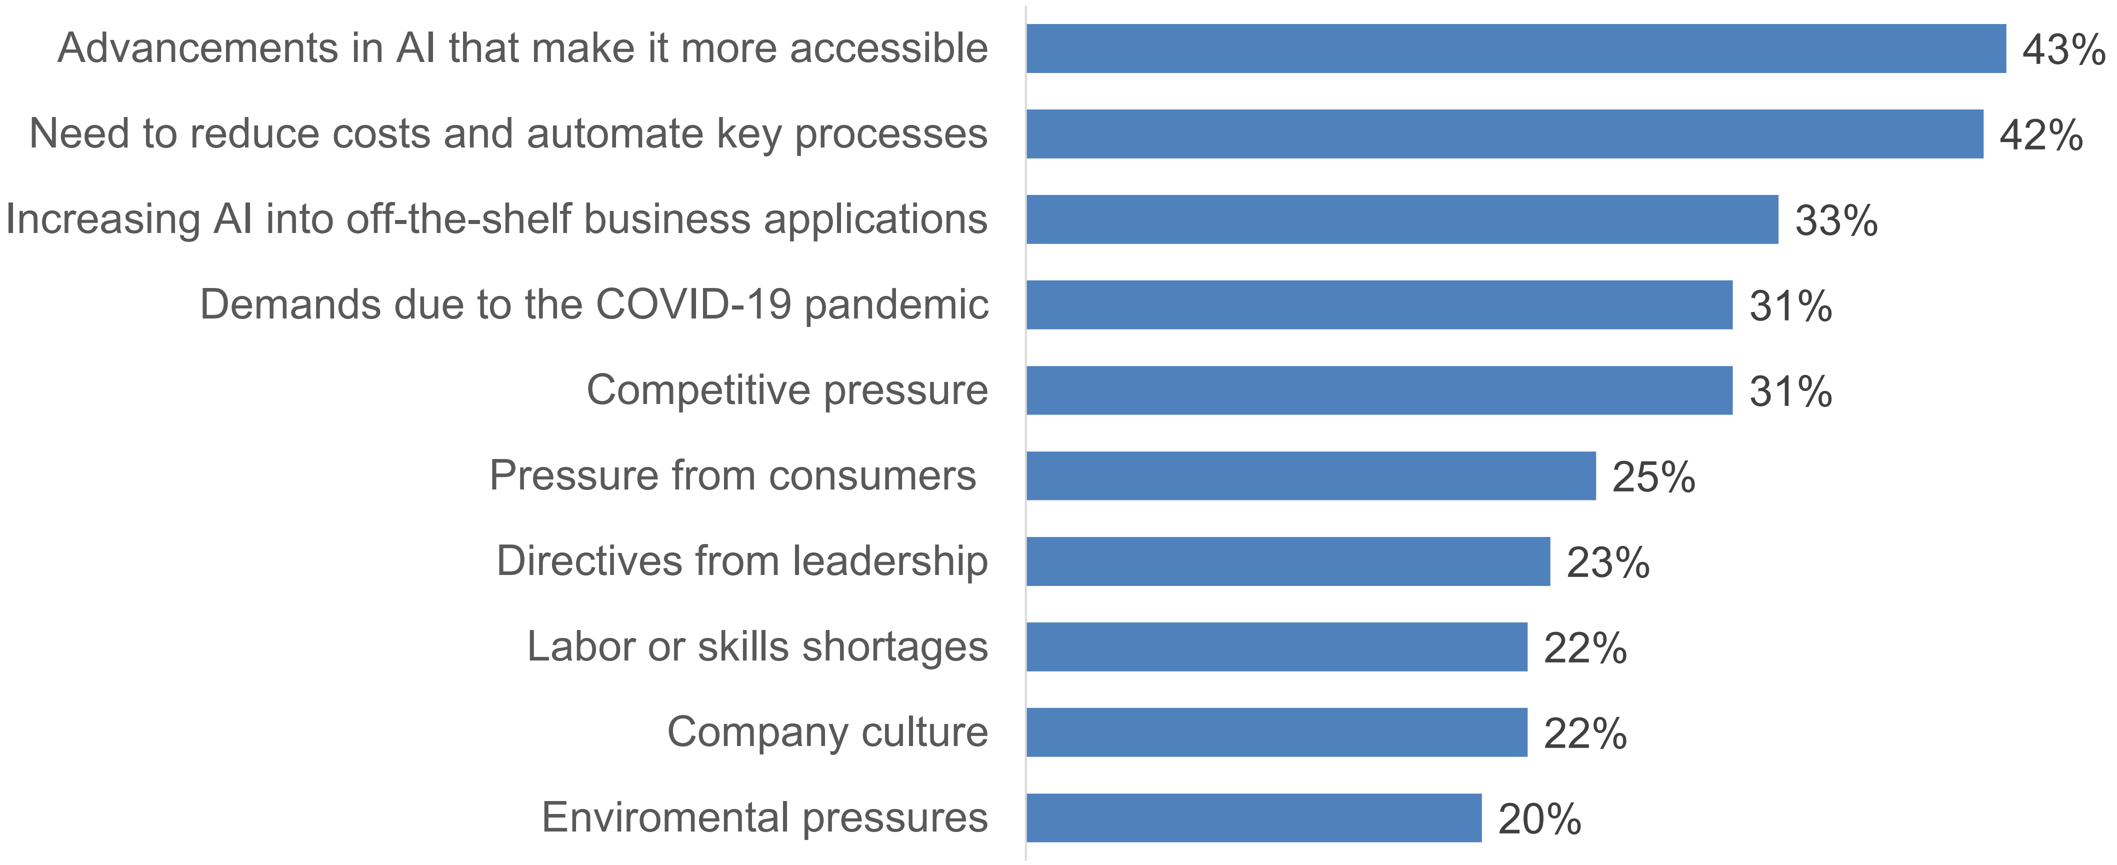

LLMs vs. Non-LLM: Clarifying the Landscape

What is the difference between traditional AI systems that have been around for 30 plus years and these newer large language models like ChatGPT?

LLMs vs. Non-LLMs

ChatGPT

General Purpose LLM

Every webpage

Books

Articles

Blogs

Social

Wikipedia

User Feedback

CoPilot

Search Model

SharePoint Files

Emails

Teams

Up to 50 files

Files up to 20 Pages

Closed LLM

Custom LLM

Staff Reports & Presentations

Master Plans

Strategic Plans

Municipal Code

Parcel-Level Data

Contracts/RFPS

Data Requirements	Traditional ML	LLMs
Data Type	Often structured/tabular (e.g., CSVs, sensor data)	Unstructured text (e.g., books, websites, code, chat logs)
Volume	MBs to GBs typically	TBs to PBs of text data
Labeling	Supervised ML often requires labeled data	Pretraining is unsupervised; fine-tuning may use labeled or instruction-following data
Diversity	Narrow domain-specific datasets	Broad, diverse datasets across many domains and languages
Data Curation	Manual or semi-automated	Requires large-scale scraping, deduplication, filtering, and toxicity checks

Privacy & Compliance	Traditional ML	LLMs
PII Risk	Lower (especially with structured data)	Higher due to web-scale scraping and potential memorization of sensitive data (ONLY FOR OPEN LLMS)
Regulatory Compliance	Easier to audit and control	Harder to trace data lineage; compliance with GDPR, HIPAA, etc., is more complex
Data Governance	Often centralized and controlled	Requires robust governance frameworks for data sourcing, retention, and redaction
Model Explainability	Easier with models like decision trees or linear regression	LLMs are black-box models, making explainability and accountability harder

Hardware	Traditional ML	LLMs
Training Hardware	CPUs or small GPU clusters	Massive GPU/TPU clusters (e.g., NVIDIA A100s, Google TPUs)
Inference Hardware	Can run on edge devices or CPUs	Requires GPUs or specialized hardware (e.g., inference accelerators)
Memory Requirements	Low to moderate (GBs)	Extremely high (hundreds of GBs to TBs for training)
Parallelism	Often single-node or simple parallelism	Requires model/data/pipeline parallelism across many nodes
Latency Tolerance	Often batch processing	Real-time or near-real-time response expected in many LLM applications

Understanding AI in the Municipal Context

How do you see AI reshaping financial operations within municipal governments?

(budgeting, forecasting, audits, contracts, or public records requests)

GFOA Collaboration

LLM of ACFRs

- Rainy day fund / reserve level as % of general fund revenue (or expenditures)
- Revenue growth rate
- Revenue per capita
- Capital spending ratio
- Expenditure per capita
- Debt per capita
- Debt per $ of assessed value or tax base
- Liability funded ratio
- Unfunded pension liabilities as % of revenue
- OPEB obligations as % of revenue

Use Cases Using Custom LLMs

- Budget Summaries
- Policy & Historical Topic Research
- Public Records Request fulfillment
- Contract Analysis
- Solicitation Drafting
- Staff Reports/Board Memos

Use Cases

Decision Support for Procurement

Pain Point:

Uncertainty around quality of raw materials procured

Solution:

Predict raw material quality based on historical data and external seasonal data

Action:

Use predictions as a part of the procurement process to make better purchasing decisions.

Value:

Make $4MM a year in additional gross margin

Use Cases

Accelerate Root Cause Analysis

Pain Point:

Mysterious process inefficiencies in a complex operation

Solution:

Identify hidden drivers of inefficiency using correlations across multiple months/years

Action:

Keep discovering drivers of inefficiency and monitor them in real time. Accelerate time to corrective action.

Value:

Make $20MM a year in additional gross margin

City of Moreno Valley

- RoadAI – PW road inventory & pavement assessment
- eSentire Cybersecurity – Log analysis from several layers of cyber defenses
- UpCodes – Building & Safety uses w/ building codes; researching for use w/ Community Enhancement & Neighborhood Services codes and ordinances
- ChatGPT – Checking inspection reports against subcontractor invoices to ensure that we are billed appropriately
- Wordly - AI to transcribe and translate public meetings; being implemented
- CitiBot - for a chatbot on our website and via cell phones
- Grammarly/MS Editor – AI to improve writing skills

High ROI Use Cases

- Detecting real-time electricity theft
- Forecasting for better purchasing decisions
- Anomaly detection for fraud & forecasting

Data Privacy, Security, and Compliance

California has some of the strongest data privacy laws in the country. How should cities manage data used by AI systems, especially if vendors and accounting data are involved?

Workforce Impacts and Job Protections

How do we balance AI-driven efficiencies with the need to protect public sector jobs and institutional knowledge?

Implementation & Success Metrics

What are the first steps to implementing AI effectively and responsibly?

Q&A

J.R. Wright

CIO @ Culver City

jr.wright@culvercity.gov

Bryant Baek

CEO & Co-Founder
@ Bettercast AI

bryant@bettercast.ai

Erica Olsen

CEO @ Madison AI

erica@madisonai.com

Steve Hargis

CIO @ City of
Moreno Valley

steveh@moval.org