Build the Expert System

Sources & Notes · Version 1.0

Canonical source: Build the Expert System - Publication Master v1

For owner-operators and accountable business or department leaders. Core companion value; free community access; no book purchase required.

Use real business evidence, keep confidential information inside approved company systems, and retain human accountability for consequential decisions.

This section is organized by chapter so the narrative can remain audio-first and free of footnotes. Historical and current-development claims are supported by the principal public sources below. Personal recollections are the author's own experience and are not listed as third-party evidence.

Current AI research, product documentation, vendor architecture, and administrative requirements can change. Those sources were checked during manuscript development in August 2026 and should receive one final publication-date freshness check before release.

Chapter 1 - Ask Alex

Federal Railroad Administration, U.S. Department of Transportation. National Train Dispatcher Safety Assessment 1987-1988, published 1990. Used for the early telegraphic train-order and dispatching mechanism and for institutional context around coordination and exceptions.

https://railroads.dot.gov/elibrary/national-train-dispatcher-safety-assessment-1987-1988

Edward Harold Mott. Between the Ocean and the Lakes: The Story of Erie, 1899/1901 edition. Used as a historical cross-check for the Charles Minot telegraph story, with caution because the account was written decades after the event.

https://upload.wikimedia.org/wikipedia/commons/3/3e/Between_the_ocean_and_the_lakes%3B_the_story_of_Erie_%28IA_betweenoceanlake00mott%29.pdf

Slason Thompson, ed. The Railway Library 1909 - Train Dispatching. Used as a counterweight to lone-inventor framing and for broader development of telegraphic dispatching.

https://www.gutenberg.org/files/50220/50220-h/50220-h.htm

National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1, 2023. Used for the broader governance point that consequential AI work requires context-specific oversight, testing, roles, and human responsibility.

https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf

Microsoft. 2025 Work Trend Index Annual Report. Used only as dated current context for organizational interest in agents and human-agent operating models, not as proof of future outcomes.

https://cdn-dynmedia-1.microsoft.com/is/content/microsoftcorp/microsoft/final/en-us/microsoft-product-and-services/ai/pdf/executive-summary-work-trend-index-annual-report.pdf

Chapter 2 - The Intern With A Vp Title

Transport for London. The Knowledge of London - Introduction to the Knowledge booklet, March 2022. Used for the qualification structure, staged examinations, and distinction between access to a map and demonstrated local competence.

https://tfl.gov.uk/cdn/static/cms/documents/introduction-to-knowledge-booklet-march-2022.pdf

Transport for London. TfL sees growth in applications for world-famous Knowledge of London during its 160th anniversary year, December 17, 2025. Used only for dated current context around the Knowledge and its training process.

https://tfl-newsroom.prgloo.com/news/tfl-seesgrowthin-applications-for-world-famous-knowledge-of-london-during-its-160th-anniversary-year

University College London / Hippocampus. London taxi drivers: A review of neurocognitive studies and an exploration of how they build their cognitive map of London, 2021/2022. Used for the sustained-learning and internalization mechanism, not as evidence about executive competence.

https://discovery.ucl.ac.uk/id/eprint/10141233/

Current Biology / PMC. Acquiring the Knowledge of London's Layout Drives Structural Brain Changes, 2011/2012. Used narrowly as supporting evidence that acquiring the Knowledge involves extended training and learning.

https://pmc.ncbi.nlm.nih.gov/articles/PMC3268356/

National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework (AI RMF 1.0), 2023. Used for the distinction between broad AI capability and readiness for a specific organizational role under defined governance and risk conditions.

https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf

Chapter 3 - Experts Live At The Edges

The Oregon Encyclopedia. Columbia River Bar Pilots. Used for the 1846 Oregon pilot-commission history, dangerous Columbia Bar context, and early licensed pilots including George Flavel.

https://www.oregonencyclopedia.org/articles/columbia_river_bar_pilots/

The Oregon Encyclopedia. George Flavel (1823[?]-1893). Used for Flavel's prior sea-captain experience, Astoria arrival, and pilot career.

https://www.oregonencyclopedia.org/articles/flavel_george_1824_1893_/

The Oregon Encyclopedia. Wreck of the General Warren. Used for the 1852 incident and the role of high-consequence local judgment, without implying that one tragedy validates every later pilot practice.

https://www.oregonencyclopedia.org/articles/general-warren/

Oregon Board of Maritime Pilots. Become a Maritime Pilot; trainee application materials; current administrative rules. Used as current evidence that professional mariner experience and supervised local exposure remain explicit parts of Oregon pilot qualification. Current requirements should be rechecked at publication.

https://www.oregon.gov/puc/bmp/pages/becoming-a-pilot.aspx

Oregon Revised Statutes, Chapter 776 - Maritime Pilotage. Used for current legal context around practical local knowledge and pilotage. Modern legal allocation is not projected backward into the nineteenth century.

https://www.oregonlegislature.gov/bills_laws/ors/ors776.html

Salesforce. AI Agents Are Smart - Handing Off To Humans Makes Them Smarter, March 24, 2025. Used as a vendor-practice example of preserving context during human handoff. It is not evidence that humans are always superior at edge cases.

https://www.salesforce.com/blog/agent-to-human-handoff/

Chapter 4 - The Company Is More Than Its Files

Hewlett-Packard Company Archives. The Sonoma Meeting. Used for the January 1957 senior-management offsite, the move toward written corporate objectives, and HP's growth beyond direct founder oversight.

https://www.hewlettpackardhistory.com/story/the-sonoma-meeting/

Hewlett-Packard. Restoring Faith in the American Economy, Carly Fiorina speech, 2002. Used for HP's retrospective interpretation that written values captured and formalized existing experience rather than appearing from nowhere in one meeting.

https://www.hp.com/hpinfo/execteam/speeches/fiorina/mit02.html

Microsoft Security. Guarding AI memory, June 22, 2026. Used as a current example of why memory that shapes future agent behavior needs lifecycle controls, auditability, provenance, and user control.

https://www.microsoft.com/en-us/security/blog/2026/06/22/guarding-ai-memory/

Microsoft Foundry / Azure AI Search. Foundry IQ: governance and enterprise AI security capabilities, June 2, 2026, and Microsoft Learn, What is Foundry IQ? Used as current vendor

examples of permission-aware retrieval, identity, sensitivity controls, and citations. They do not define the book's Company Context doctrine.

https://techcommunity.microsoft.com/blog/azure-ai-foundry-blog/foundry-iq-new-governance-and-enterprise-ai-security-capabilities/4524825

https://learn.microsoft.com/th-th/Azure/foundry/agents/concepts/what-is-foundry-iq

Chapter 5 - People Need Wins

Mary Parker Follett scholarship: International Journal of Public Administration, Mary Parker Follett's Constructive Conflict, 2007. Used for the historical distinction among domination, compromise, and integration and for the idea of identifying underlying interests rather than merely splitting the difference.

https://www.tandfonline.com/doi/full/10.1080/01900690601153148

University of Exeter. Mary Parker Follett - Rowntree Business Lectures and the Interwar British Management Movement. Used as scholarly historical context for Follett's conflict and coordination work.

https://rowntree.exeter.ac.uk/items/show/608

Dynamic Administration: The Collected Papers of Mary Parker Follett, 1941 editions. Bibliographic trail used to locate the original constructive-conflict material. The chapter paraphrases rather than quoting Follett directly.

https://openlibrary.org/books/OL14202456M/Dynamic_administration

Microsoft Research. SocialReasoning-Bench: Measuring whether AI agents act in users' best interests, May 11, 2026. Used as a controlled current research example separating task completion from competent advocacy, due diligence, private-information handling, and relationship consequences.

https://www.microsoft.com/en-us/research/blog/socialreasoning-bench-measuring-whether-ai-agents-act-in-users-best-interests/

Harvard Law School Program on Negotiation. How AI Negotiation Bots Can Deepen Classroom Learning, February 11, 2026. Used as supporting caution around incomplete information, hallucination, and negotiation guidance.

https://www.pon.harvard.edu/daily/negotiation-skills-daily/how-ai-negotiation-bots-can-deepen-classroom-learning/

Chapter 6 - The River

Toyota Motor Corporation, 75 Years of Toyota. The Birth of Jidoka. Used for Sakichi Toyoda's 1897 weft-halting mechanism and the historical point that abnormal conditions could be made visible where they occurred rather than simply propagating downstream.

https://www.toyota-global.com/company/history_of_toyota/75years/text/taking_on_the_automotive_business/chapter1/section1/item4.html

Toyota Industries Corporation. The Story of Sakichi Toyoda. Used for the longer development path, commercial trials, and loom auto-stop mechanisms. The chapter does not claim Sakichi invented the full Toyota Production System.

https://www.toyota-industries.com/company/history/toyoda_sakichi/

Toyota Commemorative Museum of Industry and Technology. Development of Loom Technology. Used as a technical-history cross-check for loom mechanization and automatic-stop features.

https://www.tcmit.org/english/research/textile/fiber03

Microsoft Learn. Multi-agent orchestration patterns and best practices; Microsoft Agent Framework handoff orchestration. Used as current examples showing that context handoff, state consistency, specialized privileges, routing, human approvals, and auditability remain explicit architecture problems.

https://learn.microsoft.com/en-us/microsoft-copilot-studio/guidance/multi-agent-patterns

https://learn.microsoft.com/en-us/agent-framework/user-guide/workflows/orchestrations/handoff

Microsoft Agent Framework. Governance at the Speed of Agents, May 14, 2026. Used as a current vendor-engineering example that orchestration alone does not remove the need for runtime policy and auditability.

https://devblogs.microsoft.com/agent-framework/governance-at-the-speed-of-agents-microsoft-agent-framework-and-agent-governance-toolkit-better-together/

Chapter 7 - Experience Bombs

National Transportation Safety Board. Trans World Airlines Flight 514 investigation, DCA75AZ005. Used for the official accident record and probable-cause framing around premature descent, procedures, and responsibility misunderstanding.

https://www.ntsb.gov/investigations/Pages/DCA75AZ005.aspx

NASA Aviation Safety Reporting System. CALLBACK Issue 435 - The ASRS Celebrates 40th Anniversary. Used for the earlier United near miss, the information-flow problem, and the later development of a voluntary, confidential, non-punitive, independent national reporting system.

https://asrs.arc.nasa.gov/publications/callback/cb_435.html

NASA. Aviation Safety Reporting System Overview. Used for the current program mechanism: voluntary reporting, confidentiality, de-identification, and learning from near misses and hazards.

https://www.nasa.gov/human-systems-integration-division/aviation-safety-reporting-system-overview/

NASA Aviation Safety Reporting System. CALLBACK Issue 555 - 50 Years of Turning Safety Reports Into Safer Skies, April 2026. Used as current continuity/history support for the ASRS program.

https://asrs.arc.nasa.gov/publications/callback/cb_555.html

NIST CAISI. Cheating On AI Agent Evaluations, December 2025; Analyzing Transcripts from AI Agent Evaluations, February 18, 2026. Used for the modern point that final outputs can hide important execution behavior and that transcript/evidence review can reveal evaluation problems.

https://www.nist.gov/caisi/cheating-ai-agent-evaluations

https://www.nist.gov/blogs/caisi-research-blog/analyzing-transcripts-ai-agent-evaluations

Chapter 8 - From Actuary To Knowledge Base

Digital Equipment Corporation, Digital at Work, 1992, archived by the Computer History Museum. Used for DEC's configuration workflow, technical editors, configuration rules of thumb, and the maintenance problem that R1/XCON addressed.

https://www.computerhistory.org/pdp-1/_media/pdf/dec.digital_at_work.1992.102630350.pdf

John McDermott. R1: A Rule-Based Configurer of Computer Systems, Artificial Intelligence 19(1), 1982. Used as a primary research source for R1/XCON's purpose, production-rule architecture, domain knowledge, constraints, and regular use by DEC manufacturing.

https://www.sciencedirect.com/science/article/abs/pii/0004370282900212

AAAI / Innovative Applications of Artificial Intelligence 1990. ReValuator - An Expert System Approach to Actuarial Valuations. Used for the actuary-to-knowledge-base transfer and for the mixed representation of expert knowledge, conventional calculations, and database case data.

https://cdn.aaai.org/IAAI/1990/IAAI90-005.pdf

Microsoft Foundry and Microsoft security/platform architecture materials, 2026. Used only as current examples showing modern enterprise AI surrounded by retrieval, identity, permissions, tools, workflow, observability, and operational controls rather than model memory alone.

https://techcommunity.microsoft.com/blog/azure-ai-foundry-blog/foundry-iq-new-governance-and-enterprise-ai-security-capabilities/4524825

https://learn.microsoft.com/th-th/Azure/foundry/agents/concepts/what-is-foundry-iq

https://techcommunity.microsoft.com/blog/microsoft-security-blog/the-microsoft-ai-and-agent-platform-%E2%80%94-the-platform-behind-intelligent-agents/4539060

Chapter 9 - Teach The System To Know When It Doesn'T Know

Stanford Computer Science. Reasoning under Varying and Uncertain Information, CS-TR-82-926, 1982. Used for historical certainty-factor rationale and the distinction between certainty factors and conditional probabilities.

https://i.stanford.edu/pub/cstr/reports/cs/tr/82/926/CS-TR-82-926.pdf

Stanford Computer Science. Uncertainty in expert systems / incomplete information, CS-TR-81-837, 1981. Used for historical treatment of missing facts, uncertain rules, incomplete information, and situations where too little is known for reliable conclusions.

https://i.stanford.edu/pub/cstr/reports/cs/tr/81/837/CS-TR-81-837.pdf

John McCarthy. Some Expert Systems Need Common Sense. Used as a historical critique of expert-system ontology and boundary limits.

https://www-formal.stanford.edu/jmc/someneed.pdf

Buchanan and Shortliffe, eds. Rule-Based Expert Systems: The MYCIN Experiments, Chapter 4, 1984. Used for MYCIN's explicit treatment of inexact and incomplete information and for the historical certainty-factor distinction.

https://people.dbmi.columbia.edu/~ehs7001/Buchanan-Shortliffe-1984/Chapter-04.pdf

Association for Computational Linguistics. The Art of Abstention: Selective Prediction and Error Regularization for Natural Language Processing, ACL 2021. Used for the durable selective-prediction mechanism in which a system may withhold an answer under uncertainty.

https://aclanthology.org/2021.acl-long.84/

Findings of ACL 2026. Abstain-R1: Calibrated Abstention and Post-Refusal Clarification via Verifiable RL; Structured Uncertainty guided Clarification for LLM Agents. Used as current peer-reviewed evidence for abstention-plus-clarification and structured asking behavior. Results remain study-specific.

https://aclanthology.org/2026.findings-acl.985/

https://aclanthology.org/2026.findings-acl.2028/

Chapter 10 - Design For Test

IEEE Computer Society. Edward B. Eichelberger profile; Thomas Williams profile; W. Wallace McDowell Award recipient record. Used for the historical development and promotion of level-sensitive scan testing and design-for-testability concepts.

https://www.computer.org/profiles/edward-eichelberger

https://www.computer.org/profiles/thomas-williams

https://www.computer.org/volunteering/awards/mcdowell

NIST. Building Evaluation Probes into Agentic AI, 2026. Used as a current research parallel to designing evaluation and traceability into agentic systems.

https://www.nist.gov/programs-projects/building-evaluation-probes-agentic-ai

NIST CAISI. Cheating On AI Agent Evaluations, December 2025. Used as a warning that a passing evaluation is meaningful only when the evaluation actually measures the intended behavior.

https://www.nist.gov/caisi/cheating-ai-agent-evaluations

NIST. Assessing Risks and Impacts of AI (ARIA): Pilot Evaluation Report, November 13, 2025. Used for multi-level evaluation including model testing, red teaming, and field testing.

https://www.nist.gov/publications/assessing-risks-and-impacts-ai-aria-pilot-evaluation-report

Microsoft Learn. Review the agent evaluation checklist, 2026. Used as a current vendor-practice example of lifecycle evaluation and regression thinking.

https://learn.microsoft.com/en-us/agents/agent-evaluation/evaluation-checklist

OpenAI Deployment Safety Hub. GPT-5.3-Codex System Card - Agent sandbox, 2026. Used as a current product-specific example that agentic capability may operate across different execution environments with different permissions, filesystem scope, and network access.

https://deploymentsafety.openai.com/gpt-5-3-codex/cybersecurity

Anthropic. Set up Claude Code, current documentation as checked August 2026. Used as a second product-specific example separating local execution context from remote model processing and provider infrastructure.

https://docs.anthropic.com/en/docs/claude-code/getting-started

Chapter 11 - Expand The Search Space

University of Cambridge. Journeys of discovery: Jocelyn Bell Burnell and pulsars. Used for Bell Burnell's first-person account of noticing an unusual signal, returning to earlier evidence, changing the recording setup, and confirming regular pulses.

https://www.cam.ac.uk/stories/journeysofdiscovery-pulsars

University of Cambridge repository. Jocelyn Bell Burnell, The measurement of radio source diameters using a diffraction method, PhD thesis, 1968. Used as primary archival context for the work in which the pulsed radio sources were discovered.

https://www.repository.cam.ac.uk/items/498e1714-0a18-4a7e-9b0d-1cba4d58a681/full

Churchill Archives Centre, Cambridge. Hunting pulsars. Used as archival corroboration for the chart record and participant description of identifying unusual recurring signals.

https://archives.chu.cam.ac.uk/collections/research-guides/misc-109/

Nature. Accelerating scientific discovery with Co-Scientist, published May 19, 2026; version of record July 1, 2026. Used for the existence and design of a multi-agent scientific AI system that generates, critiques, ranks, and refines hypotheses and for bounded reported experimental results.

https://www.nature.com/articles/s41586-026-10644-y

Google DeepMind. Conjecture Machines: AI agents and the new validation bottleneck in science, July 2026. Used as DeepMind's perspective that increased hypothesis generation may shift scarcity toward validation. It is not treated as a universal law.

https://deepmind.google/public-policy/conjecture-machines-ai-agents-and-the-new-validation-bottleneck-in-science/

Chapter 12 - The Capability Exists

U.S. National Park Service. How to Make a Gun - Springfield Armory. Used as a public-history explanation of the transition from hand fitting toward machines, gauges, and interchangeable manufacture.

https://home.nps.gov/spar/learn/historyculture/making-guns.htm

National Park Service / Raber Associates. Springfield Armory historical study - gaging and interchangeable manufacture. Used as the primary detailed mechanism source for in-process gauges, objective dimensional standards, working to gauge, continued skilled hand work, and the gradual multi-contributor path toward interchangeability.

https://www.nps.gov/parkhistory/online_books/spar/spfld_armory_history.pdf

U.S. National Park Service. The Armory and the City; Springfield Armory National Historic Site overview. Used for broader production-scale, mechanization, and institutional significance context.

https://www.nps.gov/spar/learn/historyculture/armory-and-city.htm

https://www.nps.gov/places/springfield-armory.htm

NIST CAISI. Challenges to monitoring deployed AI systems, March 6, 2026. Used for the distinction between predeployment evaluation and postdeployment monitoring under changing real-world conditions and unforeseen outputs or consequences.

https://www.nist.gov/publications/challenges-monitoring-deployed-ai-systems-center-ai-standards-and-innovation

NIST. Assessing Risks and Impacts of AI (ARIA): Pilot Evaluation Report, November 13, 2025. Used as supporting evidence for model testing, red teaming, and field testing as distinct evaluation levels.

https://www.nist.gov/publications/assessing-risks-and-impacts-ai-aria-pilot-evaluation-report

ABOUT THE AUTHOR

Alex Tran is an experienced Silicon Valley corporate professional and business owner. His background spans AI engineering, product and data environments, technical and enterprise sales, executive communication, and owner-operator businesses.

He writes for owner-operators and accountable business leaders who want AI systems to carry more repeatable work without surrendering human responsibility for consequential decisions.

His teaching focuses on practical roles, workflows, evidence, feedback loops, trust, access, and governed operating systems that businesses can actually use.