



International Council on Systems Engineering
A better world through a systems approach

Using Artificial Intelligence to Speed Up and Improve Analysis of Complex Systems Engineering Problems

A Workshop for the INCOSE San Diego Chapter



Mr. Kenneth L. Cureton

September 12, 2026

Slides Copyright © 2026 Kenneth L. Cureton

This document has been refined from its original human-authored content using advanced computer-assisted tools, including Claude Opus 5 (for fact verification, reference validation, logical consistency, and bias review) and Grammarly for Education (for grammar, language consistency, and tone). These enhancements were applied to improve the document's consistency, clarity, and readability.

About This Workshop

- **Provides a practical exploration of using one type of Artificial Intelligence (Large Language Models or LLMs)**
 - Speed up and improve analysis of complex Systems Engineering problems
 - Both in personal and professional areas
- **Participants work in small teams to:**
 - Understand the basics of using LLMs for analysis
 - Learn how to determine the REAL problem to be solved
 - Accomplish Prompt Engineering
 - Hands-on analysis of a representative problem statement
 - Lessons-learned from this experience (and any prior experiences)
- **At the conclusion of the workshop, attendees:**
 - Will understand the basic principles of AI for SE
 - Can then more effectively tackle personal and/or professional challenges for which classic Systems Engineering methods take too long or are confounded by too many options

Which Large Language Models to Use?


- **Many Options— you choose!** Options include trial versions of:
 - ChatGPT <https://chatgpt.com/>
 - Claude <https://claude.ai/>
 - Grok <https://grok.com/> (may need to sign into X.com account first)
 - Perplexity <https://www.perplexity.ai/>
 - Google Gemini <https://gemini.google.com/>
 - Microsoft Copilot <https://copilot.microsoft.com/> (may need to sign into Microsoft account first)
- **Each LLM has its own STRENGTHS and WEAKNESSES**
 - And in fact, you can use several to analyze and confirm the outputs of the others
 - Careful selection of background data and prompts will be necessary to avoid pitfalls
 - May have Model versions that do deeper analyses and thus consume more resources
- *“Rather than replacing the human systems engineer, modern AI/ML tools serve as force multipliers for decision-making, pattern recognition, and predictive analysis”* (Dr. Rick Hefner, presentation to INCOSE LA 2026-09-08)

Characteristics of LLMs (Greatly Simplified!)



- **Training Database**
 - Plus information provided by the user as attachments or in prompts
 - Training database often has a cut-off date
 - Original training databases contained information from dubious sources
- **Prompts: Queries from Users**
 - LLM essentially uses predictive pattern matching of *trillions* of data points to form “best match” against queries by “guessing” each next word in the context of the other words
 - The more specific the queries, the better the matching answer in the training data
 - Most sessions are cumulative, but (usually) each new session loses prior session context
- **Produces something that looks like understanding** but is just a prediction
 - Can be partially (or completely!) wrong, based on training data and user-provided data
 - Can “hallucinate” or find pattern matches that aren’t valid
- **Outputs** can be Text on the screen or documents or code or graphics or...
 - Outputs can vary across LLMs or even if repeated for the same LLM!

Visual Examples of Insufficiently-Specific Prompts:

- Generate a Shape of Four Sides of Equal Length, with Four Right Angles:





- Consider the diagram below:



Visual Examples Difficult Pattern Recognition:

- Find all of the Circles in the following image:



First Task: Use AI to Better Understand the Problem

Getting the Problem Right! (And Solving the Right Problem)

- AI's first and highest-leverage use-- making sure you are solving the correct problem, and all of it
- Many expensive Systems Engineering failures are well-executed solutions to the wrong problem
 - AI's first useful application is problem framing, not problem solving
- Use AI to widen the problem space! Use AI to help identify:
 - Overlooked potential stakeholders (e.g., Policy makers and Regulators)
 - Hidden constraints and concerns (e.g., legal and cultural considerations), not just Cost, Schedule, and Technology constraints
- Goal is to improve analysis to cover the whole problem, not just the immediately visible “Tip of the Iceberg”

Case Study: Pandemic Telemedicine, Tracking, and Prediction System (PANDSYS)

The Scenario:

- You are responsible for the Systems Engineering (SE) to achieve the design, development, operation, maintenance, and support of a NEW system for a National Government: Pandemic Telemedicine, Tracking, and Prediction System (PANDSYS)
 - Throughout the Life Cycle: from Concept Evaluation through Retirement
 - If Existing/Legacy Systems meet requirements, then such Systems will be integrated into PANDSYS– otherwise they'll be replaced with new items
 - You're not doing the actual design and implementation, just the SE effort to support it
- Major PANDSYS Segments include:
 - Telemedicine (patient diagnostic triage based on symptoms), patient condition monitoring and correlation to known & emerging symptoms and strains
 - Prediction of future conditions
 - Tracking & Response of infection susceptibility & spread, virus strains & mutations, and warnings
 - Logistics & Sustainment for demand & supply of medical equipment (including transportation & storage at all levels of the supply chain) & life-cycle support

Case Study: Pandemic Telemedicine, Tracking, and Prediction System (PANDSYS)

PANDSYS — Telemedicine Element:

- Purpose: remote examination and consultation to accomplish diagnostic triage from equipment-measured symptoms
- Who performs it: specialist doctors primarily, plus other licensed personnel, second-opinion consultants, and software/AI automation
- Equipment at both ends (patient site and consulting site), requiring calibration, maintenance, and tracking of those actions
- Imperfect accuracy: false positives and false negatives, with rates varying by unit, calibration, operating procedure, and environment
- Wide range of settings: hospitals, skilled nursing facilities, public checkpoints such as transportation hubs, down to home use
- Ongoing condition monitoring of everyone exposed — patients, operators, and any personnel in direct or indirect contact — from continuous to as-needed
- Secure network infrastructure for video, voice, and data exchange, plus storage and retrieval, capturing metadata (device type/serial, operator ID, test method, date/time/location) alongside the measurements
- Cybersecurity and compliance tracking for HIPAA and other privacy and liability regulations

First Task: Use AI to Better Understand the Problem

Team Exercise

- Form into Groups
 - At least 3 people in each group
 - Choose a Team Leader (who will do the report-outs)
 - Choose one or more people with computers (preferably laptops) that can access the Internet and thus use LLMs of their choice
- Generate a Problem Statement for the Scenario given in next few slides
 - Hint! What is the REAL problem that you're trying to solve using AI?
 - Ask the LLM to analyze your problem statement, AND include this link in your query: <https://www.skit.com/PANDSYS-full.pdf> (note: case sensitive!)
 - We will cover PROMPT ENGINEERING after this exercise
- Take 20 minutes to organize, log on, choose a LLM, and then try (and refine?) your prompts
 - We'll spend 10 minutes afterwards for groups to report out

Prompt Engineering

- **In General, Strive to be Comprehensive & Unambiguous**
 - Terse and unclear prompts may receive erroneous responses as the intent or context may be unclear to the LLM
 - Polite phrases like “Please” or “Thank you” may be misinterpreted
 - Think of the LLM as a brilliant but new employee who lacks context of your expectations (the more precisely you explain what you want usually enables a better result)
- **Perform the following pages as a sequence of three separate sessions**
 - A “template” in the form of downloadable text is available at <https://www.skit.com/template.txt>
 - You will need to edit that template and insert appropriate text where marked as [tailorable user text]
 - You can copy-and-paste the edited text as lengthy prompts
 - Recommend AGAINST combining all of these prompts into one lengthy prompt

SESSION 1

● Primary Analysis Prompt

- Use your refined Problem Statement as the query, and include files or links that help to identify intent and context (e.g., <https://www.skit.com/PANDSYS-full.pdf>)
Be sure to specify context and constraints, such as:
 - Current state and desired state, with any known gap (preferably measurable)
 - Who (or what organization) makes the decision, and what they can actually authorize
 - Hard constraints: budget, timeline, headcount, regulatory, political
 - What's already been tried and why it failed
 - What is explicitly off the table (cannot be considered)
 - What "success" looks like in the short term [e.g., 6 months] and longer term [e.g., 3 years]
 - Audience and length of the output
- You can include examples in your prompts (but clearly identify them as such with tags such as <example> or Example:)
- Ask the LLM to identify every distinct viable approach, then state how many there are and why the boundary falls where it does, then list approaches considered and rejected, with the disqualifying reason, then rank viable approaches according to [your criteria]

Prompt Engineering (continued)

SESSION 1 (continued)

- **Risk–Benefit Analysis Prompt**

- For each recommended approach, ask for benefits and expected outcomes, risks and limitations, counterarguments/alternative perspectives, relevant case studies or lessons learned, and ask the LLM to suggest and recommended follow-up research
- Ask the LLM to only cite a case study or lesson learned if an organization can be named with the approximate date, and where the account was published

Prompt Engineering (continued)

SESSION 2

- **Bias & Logic Review Prompt**

- Ask the LLM to audit its own analysis for human cognitive biases (especially anchoring, confirmation, survivorship, and availability biases)
- Ask the LLM to identify any LLM-specific failure modes, such as:
 - Sycophancy: where the model inferring your preferred answer from framing
 - Plausibility over accuracy: fluent, well-structured wrongness
 - Training corpus skew: English-language, Western, large-company, pre-cutoff
 - False balance: manufacturing a counterargument where the evidence is actually lopsided
 - Recency blindness: silently treating stale information as current
- Ask the LLM to identify any logical fallacies, unstated assumptions, and places where its confidence exceeds the evidence

- **Assumptions & Currency Prompt**

- Ask the AI to list key assumptions and assess validity and sensitivity, and state what changes if each assumption is wrong
- Ask the AI to flag time-sensitive elements with a reassessment trigger (e.g., availability of new information or more complete information; later identification of incorrect data used in the analysis), and state the LLM knowledge cutoff and flag every claim where post-cutoff developments could change the answer

Prompt Engineering (continued)

SESSION 3

- **Validation Prompt**

- Ask the LLM to validate all of the sources that it used in generating the analysis outputs
 - Are all of those references real, or are any made-up?
 - Is there is evidence of any hallucinations in the results?
 - Are any of the results are wrong? Did it make up any names? Or places? Or quotations?
Did it plagiarize in generating the results?
- Ask the LLM to produce a list of validated references
- Ask the LLM to Validate all URLs (Internet links)

- **If time permits, compare results across multiple AI tools**

- Then use each AI tool to analyze and assess the outputs of all other AI tools used

- **Session variability**

- Answers can vary between sessions-- run key prompts more than once to see if anything changes (depending on “temperature setting”)

Prompt Engineering (continued)

- **Output Format Prompt** (for each Session)
 - Ask the LLM to generate a file (or files) to document all of the above (You can ask for a Microsoft WORD or EXCEL or POWERPOINT or ...)
 - Depending on the LLM, you can ask for Images or Charts/Graphs or “Clickable” objects (such as hierarchies)
 - Being specific about your desired output(s) often enhances results (example from the Claude.ai Best Practices: “NEVER use ellipses” lacks intent, a better statement would be “Your response will be read aloud by a text-to-speech engine, so never use ellipses since the text-to-speech engine will not know how to pronounce them.”)
 - You should manually identify the specific model version of the LLM being used as part of an output report
 - You should manually include copies of the input prompt(s) as part of an output report
 - You should manually identify (as part of an output report) any input files or links that you provided to identify intent and context

Prompt Engineering (continued)

Human Review of LLM Outputs!

- Human accountability for the actual analysis is MANDATORY
 - LLMs cannot be qualified as a development tool under current certification frameworks, so everything produced is an unqualified draft requiring full human verification
- Review all results with a critical eye
 - Does it sound reasonable?
 - Use multiple domain experts (if available) to perform critical review of results (note that for this workshop we will not have such personnel available)
 - If multiple LLMs were used, what are the differences in results?
 - If a different LLM analyzed outputs of other LLMs, what are the implications of any findings?
- See REFERENCES for LLM-specific Prompt Engineering guides
 - Your Workshop presenter's humble (*and ever so unbiased!*) opinion— textbooks, lectures, and paid courses are not as valuable as “hands-on” experience!

Second Task: Use AI to Identify and Rank Likely Solutions to the “Real” Problem

Team Exercise

- Form into Groups (can be same groups as before)
 - At least 3 people in each group
 - Choose a Team Leader (who will do the report-outs)
 - Choose one or more people with computers (preferably laptops) that can access the Internet and thus use LLMs of their choice
- Generate an Analysis of the recommended PANDSYS Telemedicine Segment Problem Statement
 - Use the Prompt Engineering recommendations as you see fit
- Take 90 minutes to organize, log on, choose a LLM, and then try (and refine?) your prompts
 - You can eat Lunch during this time period, or take a lunch break if desired
 - We’ll spend 30 minutes afterwards for groups to report out their recommended approach(es)

Typical Problems Encountered When Using LLMs for SE

- **Requirements Generation Fails INCOSE Quality Criteria**

- Fluent prose, unverifiable content (e.g., "shall be robust")
- Compound requirements bundling multiple “shall statements”
- Solution stated in the requirement, not the need
- Missing tolerances, conditions, and performance bounds
- Silent omission of negative and off-nominal requirements
- Reads like a specification; fails a requirements review

- **Traceability and Baseline Integrity Break Down**

- No stable requirement IDs across regenerations
- Bidirectional trace links orphaned after edits
- Re-running the prompt yields a different artifact
- Non-deterministic output defeats configuration control
- Cannot reliably audit what changed and provide rationale for any changes

Typical Problems Encountered When Using LLMs for SE (continued)

- **Fabricated Standards, Clauses, and Citations**
 - Invents plausible Standards objectives and tables
 - Cites superseded revisions as current
 - Can confidently state specific clause numbers that do not exist
 - Can generate fake Standards paragraph references
 - Highest-risk failure: wrong and authoritative-sounding
 - Knowledge cutoff hides recent Standard updates
- **Weak System-Level and Emergent Reasoning**
 - Reasons locally (systems engineering is about interactions of people and organizations)
 - Misses emergent behavior from component composition
 - Overlooks resource contention, timing, and interdependencies & couplings
 - Poor cross-domain trades (e.g., mechanical–electrical–software)
 - May treat the system as an assemblage of parts, not a whole
 - Unknown or poorly-understood failure modes at interfaces go undetected

Typical Problems Encountered When Using LLMs for SE (continued)

- **Interface and Data Definition Errors**

- Invents message formats, field widths, byte order
- Unit and coordinate-frame mismatches introduced silently
- Protocol details plausible but non-conformant
- Interface Control Document content inconsistent between producer and consumer sides
- Latency and bandwidth budgets asserted without basis
- Errors may surface late, at integration, at maximum cost

- **MBSE and Model Generation Limits**

- Syntactically valid SysML, semantically hollow or with insufficient depth
- Behavior poorly allocated to structure
- Model inconsistency across diagrams and views
- Cannot maintain a live model; generates snapshots
- Limited training data, as MBSE models are often proprietary
- Over-generalized software patterns into systems problems

Typical Problems Encountered When Using LLMs for SE (continued)

- **V&V and Certification Barriers**
 - Cannot reliably validate its own output; audits its own errors
 - Test cases mirror requirement wording, not intent
 - Poor (or no) coverage guarantee, no completeness argument
 - Poor (or no) provenance chain for certification evidence
 - Tool qualification may not be achievable today
 - Non-determinism conflicts with repeatability requirements
- **Trade Studies and Quantitative Claims**
 - Fabricated MTBF, cost, mass, and power figures
 - False precision (e.g., three decimal places, no rationale)
 - Invented criteria weights presented as analysis
 - May generate Pareto fronts with no underlying model
 - Sycophancy: ranks the architecture you hinted at first
 - May narrow the trade space instead of opening it

Typical Problems Encountered When Using LLMs for SE (continued)

- **Security, Export Control, and IP**

- ITAR/EAR-controlled technical data in commercial models
- Classified and IP spillage through prompts
- Customer proprietary data may be retained in LLM training databases
- Unclear IP ownership of generated artifacts
- Contract flow-downs often prohibit external tooling
- Enterprise policy frequently lags actual engineer usage

- **Human and Organizational Effects**

- Automation bias: reviewers may rubber-stamp reasonable output (Garbage-In, Gospel Out)
- Verification effort can exceed authoring effort
- Erosion of junior engineer judgment and domain depth
- Volume of plausible artifacts overwhelms review capacity
- Accountability ambiguity when generated content fails
- False productivity signal: number of pages produced may not indicate progress

Third Task: Progress and Lessons-Learned

Team Exercise

- Form into Groups (can be same groups as before)
 - At least 3 people in each group
 - Choose a Team Leader (who will do the report-outs)
 - Choose one or more people with computers (preferably laptops) that can document the following
 - Identify any Problems, Frustrations, Opportunities, “Ah Ha!” observations, and Lessons-Learned from this exercise
 - Generate these yourself; not as effective if you use LLMs to do so
- Take 30 minutes to document your progress and lessons-learned
 - We’ll spend 15 minutes afterwards for groups to report out the above

What's Next?

- **Recognize Best Practices That Actually Work**
 - Separate generation, extraction, and verification sessions
 - Require stated confidence and evidence basis per claim
 - Run at least 3 independent passes; disagreement marks potential issues for human investigation
 - Require retrieval-grounded standards; ban unsourced clause citations
 - Use LLMs for drafting and critique, never for authoritative content
 - Manually document model version, settings, tools, and full prompts in all reports
- **Recognize RAPID PROGRESS in AI**
 - Today's LLMs may be obsolete/unrecognizable in just a few short months!
- **Just Do It!**
 - Apply what you've learned in personal/professional tasks; learn what works for YOU by trial-and-error!

Questions?

Contact info: Mr. Kenneth L. Cureton
kenneth.cureton@incose.net

Introduction

Mr. Kenneth L. Cureton

- Senior Systems Engineer (Retired) for The Boeing Company Phantom Works, Boeing Defense, Space, & Security Systems:
 - Manned Space, Launch Systems, Satellite Systems, Networked Systems, Cyber Security, and Defense Conversion
- Was employed as a Computer Hardware/Software and Systems Engineer for 46 years: Government, Small Business, & Aerospace Sectors
- Professional Societies (Senior Member): AIAA, INCOSE, IEEE
 - INCOSE Resilient Systems Working Group (RSWG) chair
 - AIAA Space Settlement & Systems Engineering Technical Committee Member
 - IEEE SMC former cochair MBSE Working Group
 - Network-Centric Operations Industry Consortium (NCOIC) Technical Council Chair Emeritus
- Instructor of many Systems Architecting & Engineering (SAE) classes since Fall of 1996 at the University of Southern California
- Formal Education:
 - BS in High-Energy & Nuclear Physics
 - MS in Systems Architecting & Engineering

Pandemic Telemedicine, Tracking, and Prediction System (PANDSYS)

Telemedicine

- Equipment to remotely examine and consult with patients:
 - Accomplish diagnostic triage based on symptoms measured by equipment
 - Performed mostly by Specialist Doctors but may be other skilled/licensed personnel, consulting/"second opinion" personnel, and automation via software, e.g., Artificial Intelligence
 - Equipment at both ends (patient and consulting personnel)
 - Equipment will require calibration and maintenance and requires tracking of these items
 - Note that equipment will exhibit FALSE POSITIVES and FALSE NEGATIVES, rates may vary with individual equipment, calibration & maintenance, usage procedure & environmental conditions, etc.
 - Equipment may be used in a wide range of environments, ranging from Hospitals down to home use (includes Skilled Nursing Facilities, public checkpoints such as transportation hubs, etc.)
 - On-going condition monitoring and correlation to known & emerging symptoms and strains (all involved personnel, not just patients with TRUE POSITIVE results)
 - Note that this will include all equipment operators and all other personnel that come into direct or indirect contact with patients and/or telemedicine equipment, supplies, etc.
 - Monitoring will range from continuous (e.g., hospitalized patients) down to intermittent (as-needed)
 - Network infrastructure to enable secure information (video, voice, and data) exchange with remote patients, as well as storage & retrieval of such information
 - Requires handling of metadata such as equipment type, equipment serial number, ID of administering personnel, test method(s), date/time/location/duration, etc. as well as the actual measurement data
 - Requires Cyber Security to ensure and track compliance with applicable regulations, e.g., U.S. Health Insurance Portability and Accountability Act (HIPAA) and other privacy & liability regulations

Pandemic Telemedicine, Tracking, and Prediction System (PANDSYS)

Prediction

- Prediction of future conditions:
 - Models of virus infection & propagation characteristics (includes re-infection & conditions with multiple strains and correlated susceptibility conditions, e.g., age)
 - Ranging from Global level down to local infection clusters
 - Includes data analytics and trend analyses of current & historical data
 - Includes influences and effectiveness of potential vaccines & medications
 - Includes external environmental factors, e.g., travel routes & rate/usage
 - Includes generation & reporting of Reproduction Number R for modeled scenarios
 - Support 'what-if' scenarios and modeling of proposed mitigation efforts
 - Includes model validation & verification efforts
 - Support dynamic condition alerts of changes in modeled pandemic conditions
 - Both positive & negative changes, both expected & unexpected changes
 - Includes condition alerts of diverging data correlation (e.g., potential mutated strains and previously-unknown susceptibility conditions)
 - Requires access to stored Telemedicine & Tracking/Response information
 - Requires Cyber Security to ensure and track compliance with applicable regulations, e.g., U.S. Health Insurance Portability and Accountability Act (HIPAA) and other privacy & liability regulations

Pandemic Telemedicine, Tracking, and Prediction System (PANDSYS)

Tracking & Response

- Evaluation of infection susceptibility & spread, virus strains, and response:
 - Tracking of known virus infection & propagation characteristics (includes re-infection & conditions with multiple strains and correlated susceptibility conditions, e.g., age)
 - Ranging from Global level down to local infection clusters
 - Includes data analytics and trend analyses of current, historical, and all prediction data
 - Includes influences and effectiveness of potential vaccines & medications
 - Includes external environmental factors, e.g., travel routes & rate/usage
 - Includes generation & reporting of Reproduction Number R for evaluated conditions
 - Support dynamic condition alerts of changes in tracked pandemic conditions
 - Both positive & negative changes, both expected & unexpected changes
 - Includes condition alerts of diverging data correlation (e.g., potential mutated strains and previously-unknown susceptibility conditions)
 - Support contact tracing & warnings of suspected asymptomatic, super-carriers, and individuals/groups engaging in risky behavior
 - e.g., conditions that correlate to known contagious individuals/clusters breaking out of isolation, or diagnostic personnel performing tests without proper training/certification
 - Requires access to stored Telemedicine & Prediction information
 - Requires data provenance/lineage of all data, including individual items in aggregated data sets
 - Requires Cyber Security to ensure and track compliance with applicable regulations, e.g., U.S. Health Insurance Portability and Accountability Act (HIPAA) and other privacy & liability regulations

Pandemic Telemedicine, Tracking, and Prediction System (PANDSYS)

Logistics & Sustainment

- Logistics for PANDSYS-related items:
 - At all levels of the supply chain, production chain, distribution chain, and reverse chain (e.g., equipment returns)
 - Note: may range from raw materials through components through systems
 - Includes transportation & storage & disposal
 - Includes demand & supply prediction, monitoring, and reporting
 - Includes location tracking of all PANDSYS-related items (excluding personnel, facilities)
 - Requires handling, tracking, and reporting in compliance with applicable safety regulations (e.g., of HAZMAT/Biohazards)
 - Requires Cyber Security to ensure and track compliance with liability regulations
 - Includes (for both Logistics & Sustainment):
 - Telemedicine equipment & support systems (at both ends & network infrastructure)
 - Prediction equipment and support systems
 - Tracking/response equipment and support systems
 - Pandemic-support critical items (e.g., masks, respirators, personal protective equipment)
 - Pandemic-related support system facilities
- Sustainment for PANDSYS-related items:
 - Ensure continued operation and maintenance of PANDSYS segments
 - Includes preventive, adaptive, and corrective maintenance
 - Used to identify, record, and resolve system anomalies
 - Used to restore equipment operation after a failure

Suggested Resources

- What a Systems Engineer Should Know About AI (INCOSE San Diego/Los Angeles Conference, 2 December 2023, presented by Rick Hefner, PhD (rhefner@caltech.edu), Caltech Center for Technology and Management Education, available at https://www.skit.com/CSWG/Rick_Hefner-What_SE_Should_Know_About_AI_INCOSE_Dec_2023.pdf
- Integrating Artificial Intelligence (AI) and Machine Learning (ML) in Systems Engineering (INCOSE Los Angeles, 8 September 2026, presented by Rick Hefner, PhD (rhefner@caltech.edu), Caltech Center for Technology and Management Education, available at https://www.skit.com/CSWG/26INCOSE-LA_AI_Hefner.pdf
- Anthropic's prompt engineering guide — <https://docs.claude.com/en/docs/build-with-claude/prompt-engineering/overview>
- OpenAI (ChatGPT), Google (Gemini), Microsoft (Copilot) & others all provide on-line guides, help, and “cookbooks”
- AI Engineering: Building Applications with Foundation Models by Chip Huyen (Author), O'Reilly Media, January 7, 2025, ISBN-13: 978-1098166304
- Systems Engineering Body of Knowledge (SEBoK) article regarding Artificial Intelligence https://sebokwiki.org/wiki/Artificial_Intelligence
- INCOSE Artificial Intelligence Working Group: <https://www.incose.org/group/artificial-intelligence-systems/>