Human-Centric Digital Engineering: Augmenting NASA Workforce via Digital Engineering

Terry R. Hill
NASA HQ Office of Chief Engineer
300 E Street SW
Washington, D.C., 20546
terry.hill-1@nasa.gov

Patricia E. Nicoli
NASA HQ Office of Chief Engineer
300 E Street SW
Washington, D.C., 20546
patricia.e.nicoli@nasa.gov

Abstract

NASA's Digital Engineering (DE) Program aims to modernize engineering practices by integrating advanced digital technologies with human creativity to deliver predictable, trustworthy, and mission-ready outcomes. This initiative emphasizes a human-centric approach, guided by principles of sensibility (Kansei) and harmony (Wa), ensuring technology augments rather than replaces human judgment. The paper explores NASA's strategic shift toward data-centric processes, lessons learned from past transformation efforts, and accomplishments such as SysML-based digital twins, generative design, and enterprise architecture deployment. It highlights the role of artificial intelligence (AI) as an assistant—not a decision-maker—while addressing ethical considerations, interoperability challenges, and cognitive load reduction to preserve human creativity. The vision for the future underscores harmonizing emerging technologies like AI, AR/VR, and digital twins with human oversight, ensuring transparency, accountability, and resilience in engineering processes. Ultimately, NASA's DE journey seeks to accelerate innovation, reduce lifecycle risk, and maintain public trust while safeguarding the irreplaceable role of human ingenuity via judgement, ethics, and values.[1]

Keywords

Digital Engineering (DE), artificial intelligence (AI), Frankenstein's Monster, SysML, Digital Thread, Enterprise Architecture, Lessons Learned.

Introduction

Humanity has long been concerned with the larger ramifications of technology on society and fundamentally who we are as a species and struggling to find harmony between the two. Some of the earliest remaining writing works in this vein date back to the ancient Greeks with the story of Prometheus which is traditionally credited to the tragedian Aeschylus (c. 525–456 BCE). [1] The Titan Prometheus, whose name translates "forethought" or "forethinker", stole fire from the gods which had been withheld from humankind by Zeus due to previous deceptions

[1] This abstract was created using a LLM to emphasize how technology can support the human in creating a concise summary while leaving the main body of work a product of human creation.

of the Titan. With the new technology of fire, humankind was then able to revolutionize their world with the development of metal working. However, this gift of technology came at a dire cost to Prometheus while seemingly leaving humanity untouched, but through the storytelling the subtext questions if some responsibilities might lie with the human beneficiary.

The question of the impacts of technology is directly revisited in Mary Shelley's 1818 novel, *Frankenstein*, appropriately subtitled as *The Modern Prometheus*, which is widely interpreted as a statement about the fear of technology and the potential consequences for humanity. Written during the Industrial Revolution, the novel serves as a warning about the dangers of unchecked scientific ambition and innovation without moral or ethical guidance.

And while technology has continued developing at an arguably geometric expansion since the days of Shelley, the largest potential technology disrupter in modern history very well could have arrived with the first public release of artificial intelligence and soon to come Artificial General Intelligence (AGI), Artificial Superintelligence (ASI), and ultimately with Self-Aware AI. These new technologies in theory could rival and ultimately surpass the capabilities of the best of humanity, thus leaving the existential question of where we and our descendants fit into such a world.

So in alignment with the themes of this conference, it is worth looking at this age-old, technology-fueled question through the lenses of the Japanese concepts of Kansei (sensibility), and Wa (harmony) so the human purpose and the irreplaceable power of human creativity remain central, visible, and a deliberate framework in how humanity charts the roles of technology, with respect to ourselves in the next decade, with far reaching effects into the future.

NASA's Digital Engineering (DE) Program is tasked to modernize the Agency's approach to engineering by blending cutting-edge digital tools with human ingenuity, judgment, and mission experience. The system under development can be defined within the Agile Systems Engineering Life Cycle Management (ASELCM) Pattern as the System of Innovation [2]. This change is being driven by several factors ranging from the increase in risk and complexity of missions being asked of NASA, to industry technical maturity which is changing what NASA is needed for, and thus the nature of associated contracts, and in turn the respective roles and responsibilities. Additionally, mission demand is accelerating the pace of change within industry (accelerated drive for innovation all-the-while reducing cost and schedule) and associated industry partners, which could question the relevance for NASA if it doesn't adjust its portfolio of capabilities and associated toolchest.

Driving NASA's DE initiative is the core principals which align well with Kansei and Wa: humans set mission/design intent and ethics; technology accelerates insight, design, iteration, and scaling. And with NASA being a civilian government agency public trust, transparency, and accountability demand human-in-the-loop oversight. This includes strong data management processes, traceability, and need for consistency, validity, reliability in order to ensure a responsible use of AI, and other future technologies going forward. This approach isn't new for NASA or government institutions, it is expected that government employees perform due diligence when executing contracts and with the introduction of AI into the environment, there will be required new thoughts put in to understand appropriate use in these functions.

NASA's journey will deliver safer solutions and capabilities by reducing risk earlier in the lifecycle, with shorter times from concept to operations through the use of technology which has implications to lifecycle costs, more resilient outcomes while amplifying human creativity.

Past Transformational Efforts & Lessons Learned

Digital transformation is not new to NASA and has been taking place from the very beginning of the Agency with the development of the computational capabilities which took computers the size of small buildings and miniaturized them such that they could fit on spacecraft and thus setting the stage for personal computers. NASA also created much of the early engineering analysis software, which was then later transitioned to the private sector, NASTRAN (NASA Structural Analysis) as an example, where originally for aerospace, it became the bedrock for Finite Element Analysis (FEA) software across multiple domains.

However, this evolution within NASA occurred organically and largely driven by immediate need vs a larger strategy and without the benefit of knowing the information technologies in the future which would define the early twenty first century: cheap computation and storage, the internet, global commercial cloud, machine learning / artificial intelligence, and a resultant data-driven global economy. Consequently, even today much of the Agency's data and information resides in un-federated information silos. [3] Additionally, the transition from paper to electronic documents in the late 1990's and early 2000's did speed up the ability to update documents and share them but still maintained the curation of critical information in documents and increased the challenge of knowing if the document in your possession was the latest information or not.

In the last thirty years, much of the transformation efforts were grassroots, usually domain-specific, and resulted in limited success or just benefiting a small portion of the Agency. In 2020, the Agency put forth one of the largest, coordinated transformation initiatives to date [4], and the engineering domain under the leadership of the Headquarters' Office of Chief Engineer (HQ OCE), made significant progress. OCE led the creation of a vision for NASA engineering and has begun laying the foundation for common digital engineering approaches (digital engineering as the seamless flow of information across the lifecycle of a project or program via processes, workflows, and connected systems), capabilities, and enterprise assets for the next decade.

> "The goal is to turn data into information, and information into insight." — Carly Fiorina, former CEO of Hewlett-Packard

The following are some key lessons learned by NASA over the last few decades of digital transformation:

- Focus on transforming core business processes to data-centricity. Organizational / business processes exist to ensure products and services are consistent, reproducible, ensure quality, safety and any other desired metric and measure. For organizations or companies which are more than two decades old, there is a very large chance these processes are based upon a document-centric approach rather than the now needed data-centric approach, and thus this is the first, and most important place to start with any organization's digital transformation. [5]

- Start with mission use cases, not technology. What problem are you trying to solve or efficiency desired to be realized? Just because we can, should we? This last question is ethically based on the use of data and information, now with the use of AI, and how either adds to the human creativity and experience or decreases the harmony between humans and our technology; to the dystopian-extreme of a minimal role for humanity.

- Co-create DE vision and architecture approach with end users and stakeholders; don't "tech-push." For example, sometimes the right solution is investment in data management vs buying new tools, and this is known by the implementing personnel, vs management or organization 'influencers' who are chasing shiny objects.

- Invest in data quality, structure, and interoperability from day one. Avoid data managed in proprietary frameworks — data management and software tools which require a lot of customization or are not largely compliant with leading interoperability standards and data exchange protocols. Keep in mind, that the need for software companies to continuously innovate for market growth means that in practice most of them will outpace the industry standards organizations and in practice will never be in full compliance, so one must prioritize needed interoperability when selecting solutions.[6]

- Keep in mind the technology and data/information storage approach today will be obsolete 'tomorrow', so always have an exit strategy such that when the inevitable upgrade / migration is required, it ideally should be relatively bloodless extract, transform, and load effort.

Digital Engineering Accomplishments to Date

NASA initiated the Digital Transformation initiative in 2020 and since that time the digital engineering component has realized many successes of which a few will be highlighted here to illustrate how a deliberate approach to transformation starting with foundational change which then allows for more transformative change [5].

Early in the transformation effort there were several pilot projects initiated to understand what was really being widely talked about in the digital transformation circles with the often-conflicting descriptions and capabilities of what was being promoted and discussed in papers, conferences, and proposals which all were lovingly referred to as "buzz word bingo." NASA internally developed the Artemis I Orion Electrical Power System (EPS) Digital Twin (2020-2022) which was based in SysML [7]. The Orion digital twin was quite successful in providing the ability to integrate SysML®, along with requirements, as flown design, and ability to incorporate live telemetry to create a true digital twin of the EPS, where SysML worked well as this digital twin was a limited and had well defined functional scope. Along with this design insight into digital twins, it provided valuable refinement of what a digital twin is and is not, of which such definitions are captured in [7] and later refined in [8] (2024-25) in the context of digital twins for failure prediction. During this time there were several pilot projects undertaken to understand how to connect tools, the growing glueware market, needed interoperability requirements, and industry standards to create an engineering digital thread.

Meaningful and lasting transformation doesn't begin with first picking new tools or even the integration of the tools and data systems in place; it begins with understanding why the organization exists and the processes and data input/output of an organization are core to providing the products and services which keep the doors open and the light on. This is discussed in detail in [5] and NASA's Engineering domain's related activities (2022-2025) in modeling the Agency-level processes and related lessons learned [9].

NASA utilized SysML models to provide direct traceability from NASA decadal exploration mission objectives to current programs, and all the way down to implementing organization's tech dev efforts (2019-present) providing the first ever linkage from strategic mission objectives all the way down to technical development and implementation efforts. [10]

Starting in 2020 the Agency began deploying commercial structural design tools which utilize machine learning to provide generative design with constraint-aware optimization to inform designers with options to maximize performance while minimizing mass. Engineers used this new technology to iterate beyond status-quo geometries to then apply a form of Wa to balance manufacturability, maintainability, and planned performance. Experience to date show a realization of 20% increase in system performance with a 30-40% reduction in mass.

In late 2025 NASA's Digital Engineering Program baselined Phase I of the DE enterprise architecture tactical plan which was crafted from input from implementing organization's DE leads (DE thought leaders), Center's IT Engineering Leads (those who often have the responsibility of deploying, maintaining, and paying for engineering tools), and finally with the Center's Engineering Directors. It should be noted that the NASA DE Program's enterprise architecture approach was shaped by mapping the needs, goals, and objects developed at part of typical system engineering activities, to known or proposed digital engineering use cases which then informed where Agency processes, information repositories, tools needed to be addressed, and was used for budget planning and task prioritization starting with fiscal year 2025.

some examples of Wa (harmony) where this new technology can do much of the legwork and allowing the human to remain the creative, decision maker. NASA DE used commercial LLMs to generate draft enterprise system "runbooks", and draft enterprise standard system operational procedures, for the new DE enterprise offerings based off industry

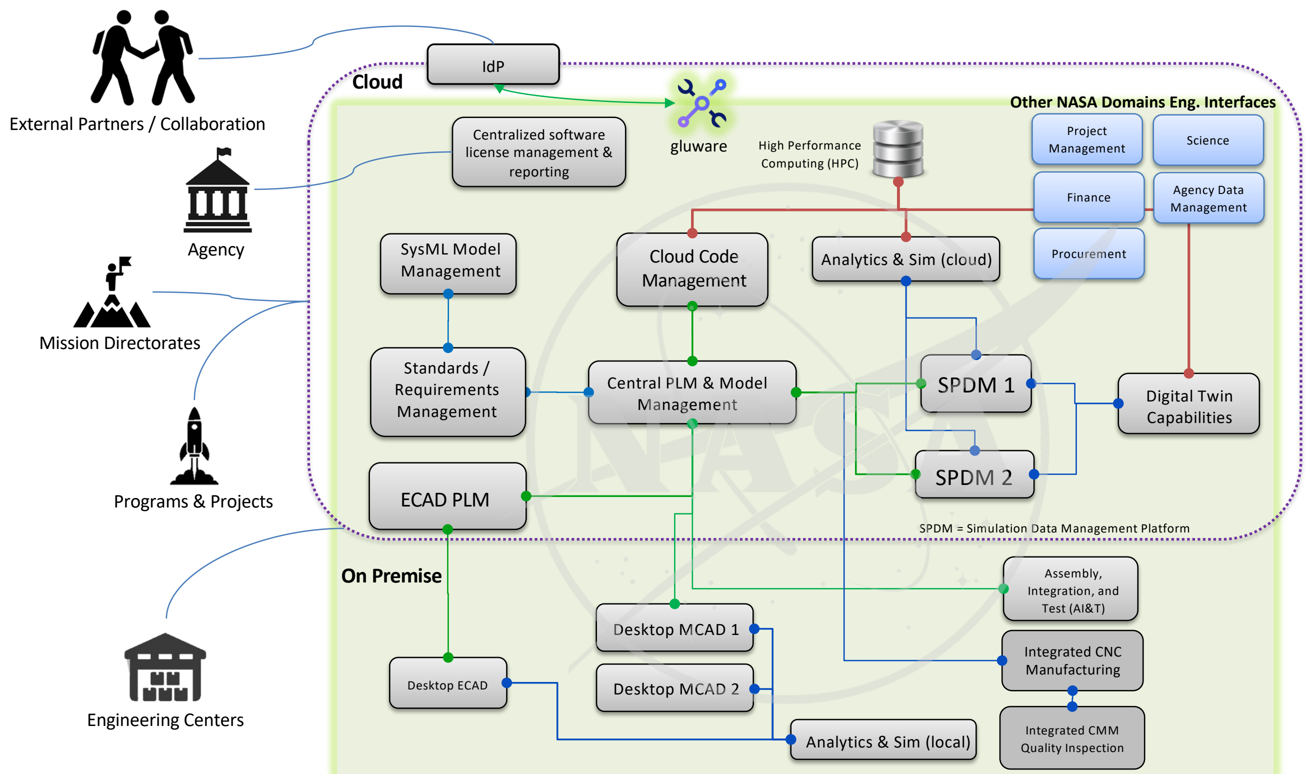

Figure 1: NASA's high-level digital engineering enterprise offering architecture with recognition that for the foreseeable future there will be a mixture of cloud-bases solutions and locally installed software but acknowledges the potential introduction of Software as a Service (SaaS).

Phase I focuses mainly of consolidation of like assets and information repositories and the federation of resultant enterprise offerings (designed/modeled utilizing Unified Architecture Framework®[UAF]), along with the license centralization of the top ten most widely used commercial software across the Engineering domain. Phase II will focus on deployment of advanced engineering enterprise capabilities which can be used by non-engineering organizations which can benefit from the DE enterprise infrastructure investment. The DE Program recognizes the "last mile" of specialty software will likely never be part of an enterprise offering.

More recently with the continued maturation of large language models (LLMs), there have been

standards and best practices. This approach saved months of labor effort, ensured consistency with similar industry approach, allowed the human to decide what was relevant to NASA DE and what was not, and even recommended some best practices not previously considered by the team and thus will improve the quality of sustainment activities. The team treated the LLM-generated products only as drafts and they have the responsibility to verify and validate.

Of course, meaningful organizational capability shift, and associated culture changes, do not happen without training and communication. To augment the existing Agency two dozen communities of practice, the DE program launched new communities of practice and DE training, related to design-to-manufacturing-and-inspection, and inter-disciplinary digital engineering (e.g. expanding Multidisciplinary Design Analysis and Optimization (MDAO) beyond just niche usage).

Planned Digital Engineering Enterprise Architecture

It is the goal of the NASA DE Program to provide a cost-sustainable, integrated enterprise digital engineering architecture (Figure 1), which includes both tools, license management, configuration/change management, and federated authoritative sources of truth, available for use by not just engineering, but any domain within NASA, all in support of realizing the NASA Mission. This enterprise offering will address GAO guidance over the last decade on centralized management of software and associated procurements, in addition to recent U.S. Government Administration emphasis on utilized commercially available solutions whenever possible, modernizing information management and cybersecurity, and increasing the efficiency and effectiveness of the government workforce.

The DE enterprise architecture will provide the engineering digital thread backbone via the federation of key authoritative sources of truth via a combination of point-to-point connectivity and glueware where necessary or advantageous.

Current Architecture Deployment Status

The deployment of the DE enterprise is underway and currently have the following elements deployed and operational:

- SysML Model Management – an Agency cloud-based, commercial solution for SysML model management. Used by programs and projects funded by all five NASA Mission Directorates to curate and share over 2.5k models.
- Product Lifecycle Management (PLM) – commercial cloud-based solution for PLM which has one NASA center currently utilizing the platform and at least three additional centers planning to migrate into it by the end of FY27.
- Glueware Server – a commercial cloud-based "glueware" solution for connecting tools and information repositories.
- High-Performance Computing (HPC) – NASA Advanced Supercomputing (NAS) Facility in addition to utilization of commercial cloud HPC services.
- Centralized Software License Management – the HQ OCE currently host cloud-based licenses servers for a SysML commercial solution and used by most of the Agency, surge pools for commercial thermal analysis and computer aided simulation and analysis software packages used across the engineering domain. The current FY26 activity will add central management of structural/mechanical computer-aided design (MCAD), electrical computer-aided design (ECAD), primary license pool for thermal analysis, in addition to top five largest software vendors used by the NASA engineering domain.

Aligning with the ASELCM SOI [2], the deployment also includes practicalities that are necessary for the workforce to use these tools, such as modeling templates, examples, best practices, handbooks, and informational sites. Pairing the deployment with updated processes and workforce support has helped to foster adoption of the technology.

Human Creativity at the Core of Engineering and Exploration: Vision for the Future

With the advent of any new technology which has the potential to cause major changes to affected industries in a short time span, it is worthwhile for organizations to spend some time visioning not only how they financially benefit but also work through any potential ethical challenges which might come along for the ride. And while replacing humans in the workplace with technological alternatives on one end of the ethical spectrum certainly stirs an emotional response, replacing human creativity on the other end might be a new area of needed examination to determine where it lies on the gradient of "even though we can, should we?" By working through different day-in-the-life scenarios of adoptions many of the concerns and hidden pitfalls can be surfaced early on and thus mitigating potential risks. The following are some engineering-related use cases presented, not to be comprehensive, but to give some introductory examples and spur your own analysis and internal discussions.

Human-in-the-loop Milestone Reviews

As programs and projects become larger and more complex the nature and intent of design reviews in some cases have become an obsolete practice where the time required to do the same due diligence given to past projects of similar size and complexity is often not given today due to pressures to move from concept to operations at an ever-shortening pace, and the risk of missing errors during said reviews is now part of the business calculus.

Moving forward the preparation for program/design reviews will change with the progression of data-centric digital engineering practices, such that with some forethought and planning the review products effectively build themselves as the program progresses, either through high-value or critical metric dashboards, reports auto-generated from design, SysML models or PLM systems – any of these from human-generated templates, or scripts, or from AI-assisted capabilities. One of the key areas to investigate reducing the time it takes for major reviews, and the voluminous amounts of information which must be consumed and understood to effectively meet the technical requirements of the review, falls into decreasing the overall cognitive load[2] which generally defined by the following three categories:

- Intrinsic load (necessary) - The inherent mental effort of the task (e.g., complex trade studies, multi-physics reasoning, system-of-system assessments, capturing lessons learned, etc.).

  Management of it requires sequencing and scaffolding of complexity, identification and validation of assumptions, it cannot be dumbed down through summarization.

- Germane load (valuable) - Effort used to build mental models and expertise (e.g., understanding rationale, patterns, capture of project knowledge, forecasting into the future, etc.).

  Encourage it with explainability, aligned ontologies, consistent data structures, and clear rationale capture.

- Extraneous load (unnecessary) - Overhead caused by poor tool user interfaces, cluttered or overloaded dashboards, information fragmented across multiple tools/repositories, inconsistent semantics, conflicting definitions or lack of common ontologies or design standards.

  This is the category of cognitive load to aggressively reduce via information management governance, common definitions and cross-domain ontologies, common design and construction standards, and organizational tool interoperability requirements via associated industry standards, and should be treated more like an organization social contract (a Wa contract if you will) rather than just compliance specifications. Organizational workflows and processes' *system telemetry* which consists of things like the number of mouse clicks required to complete a task, steps per workflow, and query

[2] Cognitive load is the total mental effort required for a person to understand information, make decisions, and complete a task.

latency, should also be examined for associated extraneous load.

Reducing unnecessary cognitive load frees engineers to focus on creative problem solving and judgment of what advanced tools like AI might present which will be the heart of harmonizing technology (Wa) with human creativity and associated out-of-the box thinking which AI is still incapable. Of course, how this reduction in cognitive load (the thinking overhead) is addressed will require a Kansei (sensible) approach reflected in and improvement of users' perceived clarity of the information presented, associated trust, and the confidence in the support of decisions.

Such 'Kansei Reviews' should address the cognitive load imposed by tools, data, and processes—beyond the essential complexity of the engineering problem itself. Data-centric management of the supporting data and information (ideally managed and curated as structured data in federated authoritative sources of truth) coupled with advanced tools, such as AI, to support the surfacing of relevant information to inform decisions in addition to combing the expansive corporate knowledge archives for relevant past lessons learned to avoid repeated undesirable outcomes.

Examples of actions to reduce extraneous cognitive load and in-line with Kansei principals:

- Cross-tool continuity: tool single sign-on, deep links across repositories (authoritative sources), and no re-entry of the same data
- Automate processes and data entry: Utilization of workflows within the engineering software tools will ensure that processes are executed correctly and data is inherited versus having to be re-typed multiple times.
- Consistent semantics: published and shared ontology, units, and naming across tools and repositories, utilized to ensure the requirements match the models and traceability between the models exists and is meaningful
- Role-based views: planner vs. analyst vs. maintainer get tailored information layouts
- Progressive disclosure: show essentials first and let experts drill down via source links
- One-glance status: clear "state of system" dashboards, tiles, and trend sparklines of critical metrics which are actionable and not vanity metrics
- Explainable uncertainty: confidence bands, "why this suggestion" justification/rationale for AI presented information, and links back to the authoritative sources
- Decision support: Inline rationale assumptions should be documented and treat decisions as managed, baselined information

This approach retains human-centric creativity and judgment as part of these critical decision gates and through addressing the extraneous cognitive load. It also restores the harmony between what we are good at and the technologies which we build with the intention or making our lives easier, and focuses the question we now face: What should the "experience" of designers/analysts/reviews be in light of modern and future technology?

AI suggests, Humans Decide

It should be stated again the goal should be that technology, including current and future versions of AI, should augment the human experience and help realize the creativity and concepts. Therefore, how do we embrace this concept without slipping down the incredibly slippery slope of deferring and delegating our responsibilities to these technologies when under schedule and financial pressure, or even regrettably, to simple apathy?

Engineers with AI assistants for error checking / optimization; human writers with AI document writing-assist, continuous AI surveillance of lessons learned throughout lifecycle of a project given new/updated system/program information and latest trending, and analysts with synthetic data simulators. All of these are examples of where AI suggests or creates and humans review and decide what to do with said information, which is critical to maintaining balance, keeping the decision-making responsibility with the human.

An example of "AI suggests, humans decide" is with some simple bracket designs done within NASA engineering a couple of years ago. The design goals were for a bracket to attach at two points with known loading, minimize the bracket mass, and avoid natural harmonics due to launch vibrations. (figure 2, table 1)

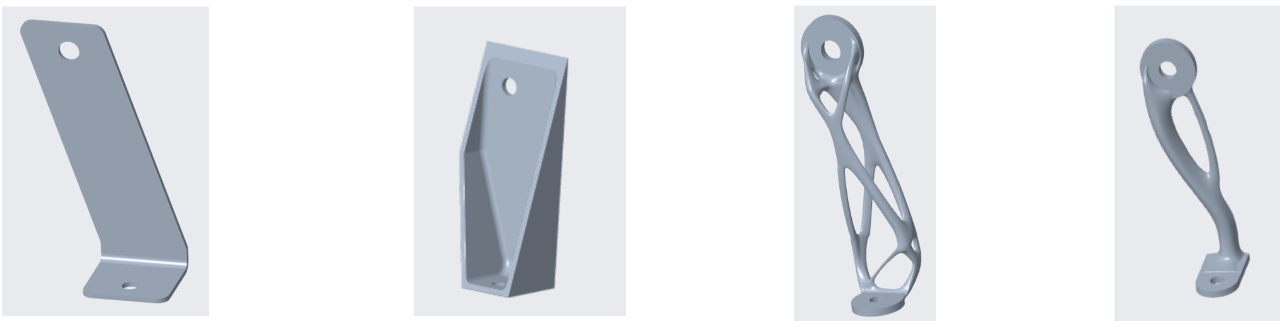

Sheet Metal | Traditional subtractive design | GTO | GTO w/ Machinable Design Considerations

Figure 2: Bracket designs using different traditional and Generative Topological Optimisation (GTO) and GTO-inspired design

Bracket	Mass (lbm)	Δm (lbm)	Mode (Hz)	Percentage of Mode Target Achieved
Sheet Metal	0.029	N/A - baseline	157	50%
Traditional subtractive design	0.080	0.051	478	157%
GTO	0.021	-0.008	513	169%
GTO w/ Machinable Design Considerations	0.030	0.001	567	187%

Table 1. Design performance between traditional design practices and GTO and GTO-inspired design.

On a similar vein, the concept of "recourse by design", when humans disagree with AI outputs, how they annotate, override, and feed corrections back into models should be talked and trained widely through organizations as an expected and intentional business practice now as the technology is still maturing and being rolled out into industry and government.

And it should be worth mentioning how technology, specifically AI, can support human engineering creativity going forward by surfacing non-obvious options which may lie undiscovered outside of current organization corporate knowledge or even human experience but fall within the known natural sciences. Managing AI-generated information or from "legacy" humanly computational and analysis practices, it is critical to integrate configuration management in into the organization business processes. Configuration Management, arguably the cornerstone of systems engineering, should be in charge of putting in place clear practices for validation, release, and pedigree of data/information, particularly important in the early days of artificial intelligence.

The notion of sensibility is also the art of balancing and mixing several points of views and systems thinking. Knowing how to ponder some elements in a situation based on lessons learned and sometimes a certain amount of undefined human intuition and critical thinking, makes the difference when addressing a complex analysis and further emphasizes the role humans should play in this new, hybrid engineering environment.

Reality, But Wait … There's More!

While much discussion has been dedicated to AI and maintaining the balance with the human, there are developing and growing areas of technology outside that realm which warrant similar discussion — particularly with new capabilities potentially giving the false sense of confidence in not maintaining the understanding and skill within the workforce.

Digital twins, currently a hot buzz phrase with a wide spectrum of varying definitions, in general is the digital representation of something which exists in reality. [8] While this technology has been around in one form or another for decades, the current groundswell of movement is toward a more holistic modeling (virtual representation) of the twinned components in the everyday physical space. And now with computational speed and data storage resource costs at an all-time low, the push toward "modeling the universe" has never been more attractive and in some cases somewhat feasible. Today there are digital twins linked to real-time operations and sustainment counterparts, which facilitate faster-than-real-time "what if" scenario testing, advanced failure / maintenance prediction, and provide a more comprehensive integration of different systems to better understand the ripple effect across all connected in ways that were not possible with past siloed models and simulations.

And with the marriage of digital twin methodologies with immersive collaboration utilizing augmented reality / virtual reality (AR/VR), not only can it be used for remote inspections, sustaining maintenance, and training, but this and the aforementioned future scenario information can be overlayed onto the user's physical reality or immersivity.

All of this is to say that with AR, VR, or digital twins and the new capabilities they bring, doesn't imply that the human workforce with those same skills should be replaced. Rather the harmony (Wa) between the two should maintained and is illustrated by the traditional "black box" metaphor which exasperates the "garbage in, garbage out". This scenario happens when the user of a tool or process cannot discern between correct or incorrect results because they do not sufficiently understand the theory, mechanics, or process required to take the input and produce the desired output. This is not a new phenomenon with AI, AR, VR or digital twins, however, with the increase in complexity of these systems, it is easy to falsely assume they are too complex for most to understand and thus people do not try, or the interface looks so convincingly real, it's assumed all the math and physics 'happening behind the curtain' are legitimate and verified. A digital twin, AR, or VR can only be useful when the system behavior is understood and reflects the behavior of that system and is underscored with P. E. Box's quote, "All models are wrong, some models are useful."

Thus, the balance between the tools and the creator of the tools must be maintained, at least at the time of the writing of this paper, as computers and associated AI can improve upon our math, computational expansion, breath of testable scenarios, and even improve upon our tools; however, they cannot yet replace human curiosity which drove the initial exploration and resulted in these tools and technologies. So, for at least the time being, this human curiosity and associated creativity must be held in balance with the appropriate usage of the resultant technology in the daily personal and workplaces.

Requirements for Ethics, Security, and Risk Management to Consider Going Forward

Likely nothing short of the personal computer and by extension the modern mobile phone and the internet have technologies ubiquitously permeated all segments of the world economy, as will artificial intelligence technology in the next half decade. As such, it will make it all the more difficult to have a uniform set of guidelines, best practices, rules or guard rails for ethical usage.

To pull again on the example of Frankenstein's monster when addressing one of the most disruptive technologies today, artificial intelligence, there are ethical issues which need to be addressed by all segments of the economy. Had Shelly's monster turned out to magically have access to all of human knowledge, Frankenstein who had already demonstrated questionable ethics for the day, could have used the endless wealth of information flowing forth from this facsimile of a human, which had no legal identity or rights, to his own benefit. He could have then used this to his personal financial benefit, claiming it all as his own without noting from whence it came. So today we have this creative conundrum to address when using AI and LLMs to create content and how and when to identify it as such to address both legal and ethical obligations. However, the rapid pace of AI development has far outstripped both areas in question, leaving many to do as best as they can in the interim.

And even aside from the aforementioned issues, there is the concern of AI model collapse which can happen if the AI hallucinations are folded back into the general knowledge repository unchecked and tagged, where the AIs then pull from this information, treating it as vetted information and creating an additive aggregation of misinformation over time. [11] One of the first stages of said model collapse "long-tail" information where uncommon, nuanced, or infrequent specialty data points get washed out of the aggregation of the AI. As AI oversamples common patterns, these unique but critical details disappear, leading to homogenized and repetitive outputs.

Focusing on the engineering domain, this is of particular concern where these unique and edge-case critical details have been hard-won by costly lessons learned, either financially or by loss of life. So as to not turn this into a dramatic Socratic dialectic, it really comes down to information and how it is managed, tagged/cited, curated, shared, and consumed.

Like content created by humans, there are generally accepted practices of citing the source of information when it is used as part of a new body of work. When using simulation data to test out designs or vet them as part of a validation or certification effort when physical testing is not an option, it is critical to know the pedigree of the simulation and associated lower-level models which created the data being used. When pulling models to be used for analysis, interface validation, contractual agreements, or even manufacturing it is critical that the state of the design be known; is this the latest released baseline, or is it the current working version?

So too must content created by AI be cited as such when in a body of work. Models and simulations created by AI must be tagged as such along with the assumptions used in their creation and general pedigree. Information aggregated by the various AI models must be documented as such so that the end consumer of the information can treat the information with the appropriate level of legitimacy and vet it accordingly. Otherwise, one might one day find themselves in an embarrassing situation where their products are found to be riddled with AI hallucinations [12] – at best – or could result in loss of life at worst if not scrutinized properly in critical situations.

There are industry groups which are actively working these issues, and it is incumbent of engineering and other domains to keep abreast of these evolving best practices and actively roll them into the internal processes. One such groups is the Coalition for Content Provenance and Authenticity, or C2PA, which

> *"…provides an open technical standard for publishers, creators and consumers to establish the origin and edits of digital content. It's called Content Credentials, and it ensures content complies with standards as the digital ecosystem evolves." [13]*

The development of these increasingly important guardrails of transparency, traceability, have close ties to safety and quality processes maintained by humans, and are important to keep focus on in the years to come as the pace of technological change is predicted to continue to accelerate.

Summary

The themes presented in this paper regarding the harmony / balance between the roles of humans and our creations are not new but do warrant continual examination as with the advent of every new advancement we must ask ourselves, "Just because we can, should we?" to ensure our cumulative best interests are kept in the forefront.

Ethics shouldn't just be guardrails and should be an active design dimension that dignifies human contribution and creativity. This will require some cultural shifts in our Western relationship with technology and personal gain and take on more of the respect typically given to technology which can be quite dangerous if used inappropriately; not out of fear, but out of respect for the technology and the resultant world in which we live.

In the end, humans provide the mission, creativity, ethics, and final decisions and associated responsibility regardless of the technology involved. But like the basis of every story ever created by people, the bottom-line consideration is the human experience and thus it is our responsibility to dictate where we fit into a world where our own technological creations change it daily.

References

[1] Britanica. (2025). *Prometheus*. https://www.britannica.com/topic/Prometheus-Greek-god

[2] Schindel, Bill, Dove, Rick (2016). *Introduction to the Agile Systems Engineering Life Cycle MBSE Pattern.* INCOSE IS.

[3] National Aeronautics and Space Administration. *NASA Data Strategy FY2025 – FY2027*. (2025). https://www.nasa.gov/wp-content/uploads/2025/07/nasa-data-strategy.pdf?emrc=488988

[4] National Aeronautics and Space Administration. *NASA's Digital Transformation Strategic Framework & Implementation Approach*. (2022). NASA/TM–20220018538. NASA. https://ntrs.nasa.gov/api/citations/20220018538/downloads/2022-1206%20NASA%20TM%20-%20DT%20Strategic%20Framework%20%2B%20Implementation%20Plan%20(Marlowe%2C%20Haymes).pdf

[5] Hill, Terry R., Richardson, David H. (2025). *NASA's Four Layer Cake: Systematic Approach for Developing a Digital Transformation Tactical Plan for Organizations*, AIAA SciTech.

[6] Aravindhan, Manigandan, *Breaking Down Data Silos: How AI 'Builds Bridges' in the Cloud*, European Journal of Computer Science and Information Technology, July 02, 2025, https://doi.org/10.37745/ejcsit.2013/vol13n48177187

[7] Pierce, Gregory J., Heeren, Joshua D., Hill, Terry R. (2023). *Orion SysML Model, Digital Twin, and Lessons Learned for Artemis I*. INCOSE IS.

[8] Inter-Agency Working Group on Engineering Complex Systems, Targeted Action Group Report for FY24 – 25. (2025). *Digital Twins for Failure Prediction White Paper*, NASA/TM−20250008957, Johnson, Theodore (NNSA/DoE), Hill, Terry R. (NASA), Cornford, Steven (The Aerospace Corporation – NASA), Adams, Mark (DHS), Biloiu, Ioana (NNSA/DoE), Foote, Andrew (DoD), NASA.

[9] Hill, Terry R., Nicoli, Patricia E., Cornford, Steven L., Morgenstern, Robert M., Chiras, David M., Barnes, Patrick D., Bendig, Josh L. (2024). *Modeling NASA's Procedural Requirement Processes - Implications for Digital Future*, INCOSE IS.

[10] Hill, Terry R., Carnevale, Alanna E., Morris-Eckart, A., Sundaram, V., Farhaj, Leon Z. (2024). *NASA's Use of MBSE and SysML Modeling to Architect the Future of Human Exploration*, INCOSE IS.

[11] John Burden. (2025). *Model Collapse and the Right to Uncontaminated Human-Generated Data*. Harvard. https://jolt.law.harvard.edu/digest/model-collapse-and-the-right-to-uncontaminated-human-generated-data#:~:text=Model%20collapse%2C%20whereby%20LLMs%20and,produced%20by%20generative%20AIs%20themselves

[12] Weiser, B. (2023, May 27). *Here's what happens when your lawyer uses ChatGPT*. The New York Times. https://www.nytimes.com/2023/05/27/nyregion/avianca-airline-lawsuit-chatgpt.html

[13] *Advancing digital content transparency and authenticity*. (2025). Coalition for Content Provenance and Authenticity. https://c2pa.org

Biography

Terry R. Hill

Received his BS and MS in Aerospace Engineering and Guidance, Navigation and Control theory (emphasizing the use of banks of neural networks to optimize state vectors) from the University of Texas in Austin.

Terry currently serves as the NASA Digital Engineering Program Manager out of headquarters Office of the Chief Engineer (OCE) and responsible for providing an executable strategic and implementation approach for delivering digital engineering and MBSE methodology and interoperable toolchains.

His previously held positions: X-38 GN&C navigation algorithm development, ISS GN&C flight software validation, Orbital Space Plane (OSP) systems engineering and integration, leading portions of the Shuttle Return-to-Flight tile repair, as the next-gen Constellation Spacesuit Engineering Project Manager; ISS EMU (white spacesuit) Deputy Subsystem Manager; and Deputy Program Manager of the NASA Crew Health and Safety (CHS) astronaut medical program.

Patricia E. Nicoli

Received her BS in Electrical Engineering from the University of North Dakota and her MS in Industrial Engineering with a Systems Engineering Focus from the University of Florida. Patricia currently serving as the NASA Digital Engineering Deputy Program Manager out of headquarters Office of the Chief Engineer (OCE).

She also serves as the lead of the Agency's Model-Based Systems Engineering Leadership Team. As part of this function, she coordinates across the different NASA Centers with MBSE leaders and practitioners to increase the model-based capabilities of the workforce and migrate to data driven processes using modern applications and tools. The goal is to provide improved processes to enable the workforce to deliver quality products consistently and efficiently.