REPORT

Applied Artificial Intelligence and Machine Learning

Test and Evaluation Program for

Consumer Products

August 24, 2022

For further information, contact:

Nevin J. Taylor Chief Technologist

Office of Hazard Identification and Reduction

(301) 509-0264 email; ntaylor@cpsc.gov

EXECUTIVE SUMMARY

The Consumer Product Safety Act (CPSA), [1] enacted in 1972, established the U.S. Consumer Product Safety Commission (CPSC). The CPSA and other laws administered by CPSC provide the agency the basic authorities to develop mandatory product safety standards and, if necessary, ban products that present unreasonable hazards to consumers. The CPSA was amended in 2008, to provide CPSC significant additional regulatory and enforcement tools to carry out its mission. These amendments addressed areas, including third party testing and certification, civil and criminal penalties, and created the saferproducts.gov website. The CPSA charges the Commission with protecting the public from unreasonable risks of serious injury or death from thousands of types of consumer products under its jurisdiction.

In 2020, the Chief Technologist (CT) was tasked with identifying Artificial Intelligence (AI) and Machine Learning (ML) technologies in consumer products and creating a program to assess safety-related issues. In 2021, the CT was charged with creating a test and evaluation (TE) program to analyze hazards associated with AI/ML in consumer products. CPSC staff hosted two forums with industry, academia, and government agencies. The 2021 forum[2] focused on developing the ability to screen and assess AI/ML capabilities in consumer products,[3] and the 2022 forum[4] focused on analysis of products through a TE program to determine if AI/ML technologies in consumer products pose an unreasonable risk of harm for consumers.[5]

The initial AI/ML report released in May 2021, *Artificial Intelligence and Machine Learning in Consumer Products*, outlined the initial stages to screen for the components of AI/ML in consumer products.[6] The CT's report, *Applied Artificial Intelligence and Machine Learning Test and Evaluation Program for Consumer Products*, builds on the concepts from the 2021 report, describing in greater detail a proposed process to evaluate AI/ML technologies in consumer products to determine if they pose a potential hazard. The resulting four-step approach to this process: (1) screen for AI/ML components; (2) assess the contribution they make to the product; (3) analyze if they contribute to known hazards; and (4) monitor/measure their transformations. This is done through an iterative approach which monitors and measures if the products evolve beyond their initial parameters. The four steps are listed below:

- **SCREEN** *consumer products and identify the existence of AI/ML technologies*
- **ASSESS** *capabilities and determine the implications of these technologies*
- **ANALYZE** *contributing factors that AI/ML have to discern if hazardous*
- **MONITOR/MEASURE** *conditions to determine if/when AI/ML evolves beyond safe parameters*

[1] *Consumer Product Safety Act (CPSA) | CPSC.gov*
[2] *Federal Register :: CPSC Artificial Intelligence Forum*
Federal Register :: CPSC Artificial Intelligence Forum
gov/s3fs-public/2021-03-02-CPSC-Artifiical-Intelligence-and-Machine-Learning-Forum.pdf?0BasTb2v6e_70a.6rlW7izqMDTd0RTru
[4] *Federal Register :: CPSC Artificial Intelligence and Machine Learning Test and Evaluation Forum*
[5] *https://www.cpsc.gov/s3fs-public/2021-03-02-CPSC-Artifiical-Intelligence-and-Machine-Learning-Forum.pdf?0BasTb2v6e_70a.6rlW7izqMDTd0RTru*
[6] *https://www.cpsc.gov/s3fs-public/Artificial%20Intelligence%20and%20Machine%20Learning%20In%20Consumer%20Products.pdf*

CPSC is working to establish testing protocols and methodologies to analyze emerging technologies as they are integrated into consumer products to ensure practices are in place that consistently meet applicable safety requirements and specifications. One such technology is AI/ML, given its integration into consumer products. The CPSC recognizes the need to examine the application of these technologies to determine if they possess the potential to cause injury or death to consumers. This report provides an overview of an applied AI/ML test and evaluation (TE) program designed to determine if these technologies are safe as they exist in consumer products.

As noted by the presenters in the 2022 Applied AI/ML TE Forum, this is no small undertaking. Given the complexity involved, developing an AI/ML TE capability will be undertaken in three phases. The first phase will focus testing specific AI/ML technologies in consumer products. This phase will outline the requisite talent, facilities, equipment, and services necessary to evaluate a specific portion of products with AI/ML technologies posing a risk to consumers. Phase two will explore how these capabilities operate in consumer products and will look at their integration with other capabilities in an effort to evaluate their potential contribution to hazardous conditions. Finally, phase three looks at the inputs that influence the AI/ML system, such as data, how the system interprets the resulting information, and the implications of the outputs it produces. Additional considerations in this final step are the potential impacts that occur due to different environmental conditions and the unforeseen uses that could cause unexpected outcomes.

This report proposes a way forward and presents the necessary steps to establish an applied AI/ML TE program and provides recommendations for the Commission's consideration for their upcoming Fiscal Year 2023 Operating Plan. It concludes by identifying the necessary next steps to implement this program at CPSC.

TABLE OF CONTENTS

INTRODUCTION

This report presents a proposal for implementing an applied AI/ML TE program at the CPSC. The goal of this program is to support the CPSC's mission of saving lives and keeping consumers safe by reducing the unreasonable risk of injuries and deaths associated with consumer products specifically regarding those products that contain AI/ML technologies. Staff provided a 2021 report, which introduced the concepts of AI/ML and outlined federal efforts underway, relative to these technologies, and how they would apply to the CPSC.[7] That report defined "AI," described the application of ML, and talked about the need to evaluate AI/ML technologies in products to ensure they are safe for the consumer. In the report, staff identified six objectives, listed below. This report focuses on developing tests and evaluating the capabilities needed to support objectives 3, 5, and 6.

1. ***Conducting research on potential product hazards from AI/ML in consumer products.***
2. ***Identifying incidents where consumers are injured due to AI/ML in consumer products.***
3. ***Developing new or enhancing existing voluntary consensus standards with industry to address AI/ML safety hazards in consumer products.***
4. ***Informing and educating consumers on AI/ML safety hazards in consumer products through the media, private organizations, state. and local governments.***
5. ***Obtaining the recall of consumer products that pose a substantial product hazard and arranging for their repair, replacement, or a refund; and***
6. ***Issuing and enforcing mandatory standards to address consumer product safety hazards due to products with AI/ML technologies.***

BACKGROUND

The operation of consumer products has become less transparent with the introduction of AI/ML technologies. The increasing dependence on these technologies, and the functions and features they provide, has brought to the forefront the need to determine their potential contributions to hazards. To that end, CPSC is developing a TE program to analyze the potential risk these technologies could pose to consumers, and CPSC has adopted the National Institute of Standards and Technology (*NIST*) definition for "AI":

> ***AI technologies and systems include software and/or hardware that can learn to solve complex problems, make predictions or solve tasks that require human-like sensing (such as vision, speech and touch), perception, cognition, planning, learning, communication or physical action.*** [8]

For this report, "ML" is described as:

> ***"The field of study that gives computers the ability to learn without being explicitly programmed."***[9]

[7] https://www.cpsc.gov/s3fs-public/Artificial%20Intelligence%20and%20Machine%20Learning%20In%20Consumer%20Products.pdf?qtcLqPHYa9bDkWTXCgH1rMsT3QineKXQ

[8] *https://www.nist.gov/news-events/news/2019/07/nist-releases-draft-plan-federal-engagement-ai-standards-development*

[9] *Study on the Impact of Artificial Intelligence on Product Safety (publishing.service.gov.uk)*

The 2021 report established six overarching goals to ensure safe AI/ML in consumer products. It delineated those goals into four objectives for the organization that will be incorporated into an actionable applied AI/ML TE program that answers the following questions:

- **SCREEN**: *Do consumer products have AI or ML?*
- **ASSESS**: *What functional features do they provide?*
- **ANALYZE**: *How does it impact consumer safety?*
- **MONITOR/MEASURE**: *Does AI and/or ML in a product create an unreasonable risk of injury or a substantial product hazard at some point throughout its lifecycle?*

In fiscal year 2022, CPSC hosted a forum with industry, academia, and consumer groups focused on these goals, seeking to explore the best way to test and evaluate these technologies. [10] The information gleaned from this event, coupled with that captured at NIST's AI/ML working group and workshops,[11] combined with the work of the CPSC AI/ML WG, is the focus on this report. This report will take the lessons learned and best practices from this event and incorporate them into a proposed program to fulfill the CPSC's mission relative to AI/ML technologies.

At the Applied AI/ML TE forum, the National Institute of Standards and Technologies (*NIST*) highlighted the importance of identifying the building blocks of technical requirements. They illustrated the concepts, terminology, and taxonomy upon which to establish requisite metrics, evaluations, and benchmarks essential to determining the risk management framework (*RMF*) for AI/ML technologies. They highlighted their ongoing work to establish an RMF to determine if these technologies perform as specified. It is within these evolving frameworks that CPSC will determine if AI/ML presents an unreasonable risk to consumers.

An important consideration for the RMF is the data that feed the AI/ML algorithms. At the 2022 forum, the SAS Institute highlighted the importance and impact of data, identified the dependencies, and illustrated the challenges associated with an exponential growth in data, where only 10-20 percent of all data is structured, and less than .5 percent is analyzed. This has a dramatic impact on the algorithms that are fed by the increased integration of IOT devices that are collecting increasing amounts of data. With most computational capability being conducted in the cloud, it has become increasingly difficult to evaluate products, given their new roles as peripherals *(input/output devices*) for these technologies.

Due to the integration and interdependencies of these technologies, it is an extremely challenging and a complex task to test and evaluate AI/ML technologies in consumer products for the purposes of determining if they are safe. Through the application of a use-case approach, the National Artificial Intelligence Institute (*NAII*) illustrated the benefits that they observed and the impact that AI/ML capabilities were making for veterans and consumers. They shared their lessons learned and evolving best practices, which clearly characterized the contributions of these capabilities. [12]

[10] *2022_03_31CPSCMeetingLogAI_ML_TE_Forum.pdf*
[11] *AI Measurement and Evaluation Workshop | NIST*
[12] *2022_03_31CPSCMeetingLogAI_ML_TE_Forum.pdf*

At the forum, Oracle provided keen insights relative to the computational capability necessary for these technologies to exist. Unlike historical on-premises solutions, they highlighted the cloud's computational capability, given the adaptive nature, ability to scale, and accessibility to data that serve as the new standard for High Performance Computing (*HPC*). Given its adaptive capacity, flexibility, and performance, the resulting increased capability, coupled with lower cost, is contributing to the exponential growth of consumer products with AI/ML. The distributed nature of this cloud-based solution increases the complexity of TE and makes it necessary to differentiate the dependencies of products that rely on these services. Oracle emphasized the importance of distinguishing the technologies as they exist in products from the services they provide, when considering their contributions to known hazards that create safety concerns.

This was demonstrated in the MITRE presentation on micro mobility, which illustrated many AI/ML capabilities in both products and services. The distributed nature of micro mobility and its reliance on IOT for connectivity highlight growing data dependencies. This, coupled with the integration of capabilities like geo-location, is essential for unlocking and monitoring/measuring performance. MITRE's micro-mobility use case demonstrated the complexity within the systems, given reliance on AI/ML to provide capabilities to the consumer. It became clear through the course of MITRE's presentation that standards for integration, communication, and collaboration are key to the viability of this product.

Especially worth noting was Underwriters Laboratory's (*UL*) work to address standards for automation. Their UL 4600 Standard for Safety for the Evaluation of Autonomous Products is just one effort underway that addresses the impact of AI/ML capabilities that create independent products, like robots to robotic control arms. Even though current efforts to create standards are in the formulative stages, organizations like UL are addressing the growing need to provide parameters for products with AI/ML. Ongoing TE efforts will allow stakeholders to identify and inform future standards that address potential hazards for consumers. UL's current focus on integrated systems within a product highlights the need to validate and explore the potential of these technologies:

- **Argument Sufficiency** – *Identify arguments and support claims*
- **Evidence Sufficiency** – *Claims are clear and concise*
- **Validity of Evidence** – *Identify risk and criteria*
- **Accepted Risk** – *Demonstrate through lifecycle*
- **Safety culture** – *Reproduceable element*

Through validation testing of existing safety mishaps of products that contain AI/ML, the potential and probable impact to the consumer can be identified. By measuring the reliability and resulting dependency of products through a standard risk management framework (RMF), the impact from the technologies can be measured. One such example of this is UL 4600, which illustrates the safety considerations of AI/ML devices operating in an autonomous mode. In their

use case, a consumer robot is examined to consider the relative risk. Through this preliminary analysis, UL is able to explore insight on potential next steps, based on actions taken measured against relative results achieved. The following offers a flowchart of this process:

Bureau Veritas (*BV*) provided practical insight into the application of an RMF through the TE

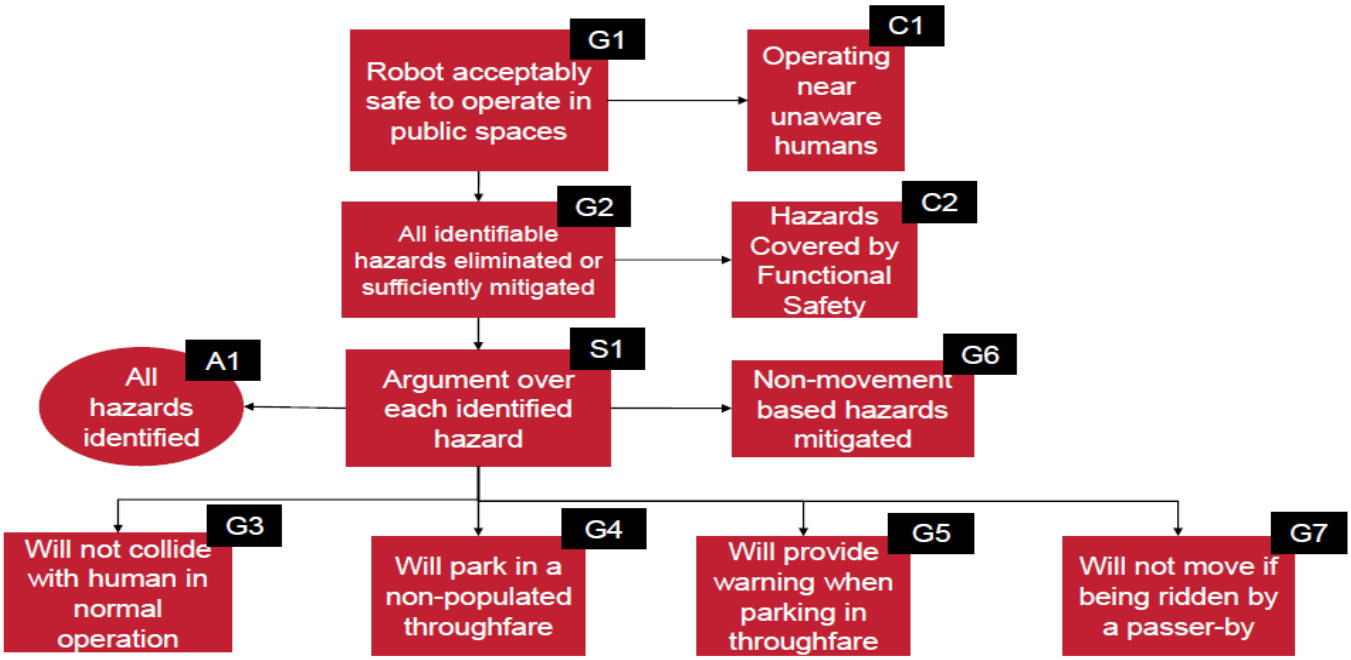

process. BV spoke on analyzing the performance of the product to determine if it is within an acceptable safe operating envelope, given current considerations. They identified the Testing Inspection Certification (*TIC*) process from the F15.75 council as an example of evaluating the AI/ML design, development, and deployment lifecycle. They presented an RMF that examined the potential of injury, given the probability of occurrence, severity of hazards, and cumulative impact. The following chart illustrated how they monitored and measured hazards that they identified and their potential impact on consumer safety.

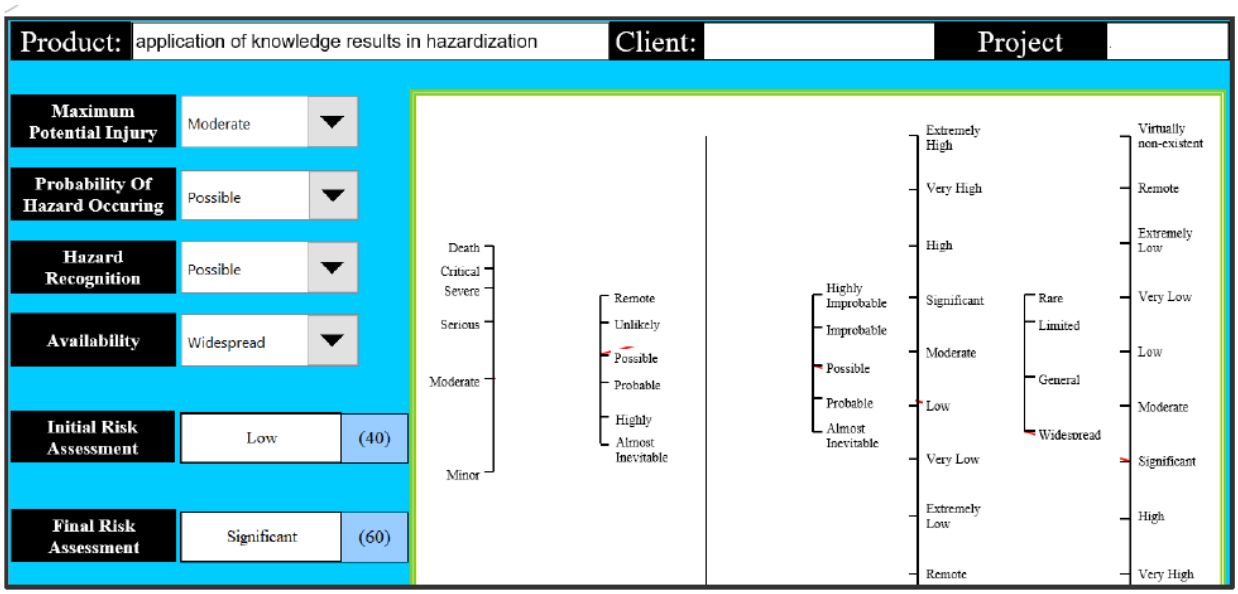

From this, a testing plan can be constructed from which to evaluate the contributing factors of AI/ML to the product's known safety considerations. It is within this initial analysis that the implication from AI/ML and its impact to hazards can be evaluated. Given the complexity of this effort, the challenges associated with identifying and addressing all the components as they are integrated into the product, and evaluating the collective whole, is daunting. It is only at that point that the implications and impact of AI/ML contributions to known hazards can be assessed.

PROPOSED TE PROGRAM

The Applied AI ML TE forum highlighted the value of a use-case approach to inform the testing and evaluation of AI/ML. With only a few programs existing to TE AI/ML, it is important to take a progressive approach to screen, assess, analyze, monitor, and measure these technologies. This process allows for the requisite determinations whether the product is evolving beyond parameters of what has been previously deemed to be a safe or if it needs to be re-evaluated. This would require a new assessment and analysis to determine through another TE effort if the product poses unreasonable hazards due to its transformation. By establishing a procedure to monitor and a process to measure AI/ML products, we develop the means to determine the potential for these technologies to create unreasonable risk to consumers.

To evaluate a product where AI/ML are present, a Subject Matter Expert (*SME*) would determine what capabilities exist, and document them into functional categories for further analysis. This uses the aforementioned four-step process, outlined below, to addresses the original four questions addressed in the first CPSC AI ML Forum[13]:

SCREEN: *identify if the product contains AI/ML components*

ASSESS: *implications of products with AI/ML capabilities*

ANALYZE: *impact of the technology to contribute to safety concerns*

MONITOR/MEASURE: *for unsafe conditions as AI/ML transform the product*

STEP ONE-SCREEN: For products that could injure consumers and may have AI or ML, the first step is to determine if the product is or could be AI capable. To screen for AI, identify if there are data that feed the system, determine if algorithms inform the cognitive process, whether connectivity exists to communicate, and whether the AI has processes that correlate knowledge into understanding.[14,15] If these components, or the ingredients required to create AI capabilities, are not present, then we can determine that there is no AI in the product.

Screening for ML is a process to determine if the system has the means to monitor internal activities. Also required is the ability to measure these activities in a way to model how they influence and affect the system. It is through ongoing monitoring of past actions and current activities that future potential hazards that create safety concerns can be identified. Thus, screening for ML can include monitoring to compare performance over time, measuring the changes, and modeling the system's behavior.

Examining a consumer robot demonstrates how this screening process works. The following is a flow chart designed to facilitate the screening process that illustrates the considerations taken to determine if the necessary AI components exist. Through this process we can identify if this robot is an AI-capable product.

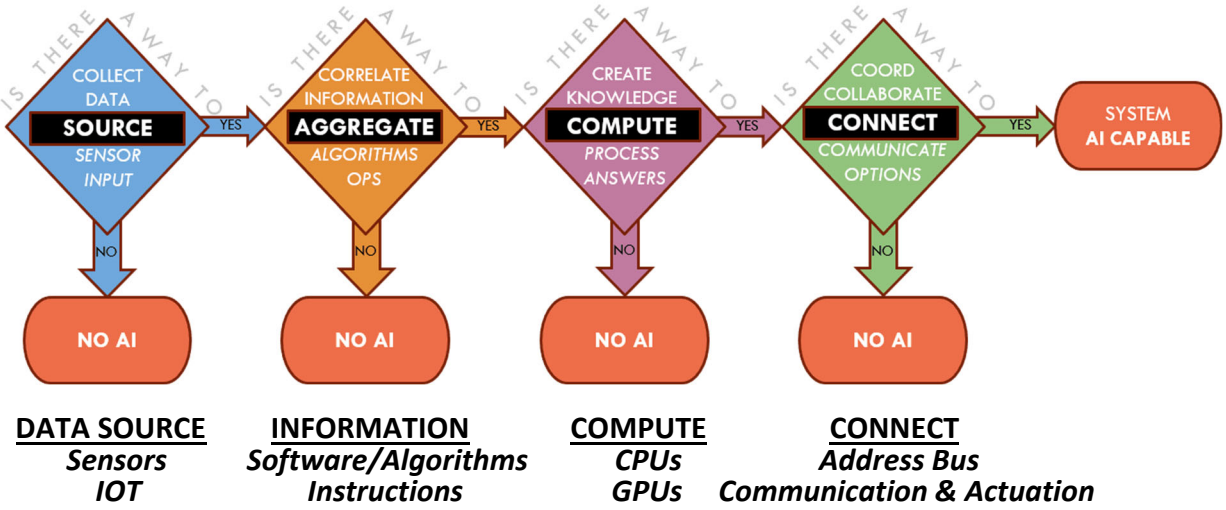

[13] Federal Register :: CPSC Artificial Intelligence Forum
[14] What to Expect From Artificial Intelligence (mit.edu)
[15] What to Expect From Artificial Intelligence (mit.edu)

The robot possesses optics, microphones, light differentiators, and proximity sensors from which to determine the conditions as they exist around it. This allows the robot to become aware of its surroundings and collect inputs from which to understand the available options and opportunities to carry out its objectives. By monitoring and measuring visual acuity, day light, audible inputs, and distance, the robot determines the best course of action. The following flow chart provides the building blocks upon which the robot learns, and thus, they are the basic components necessary for ML capabilities to exist:

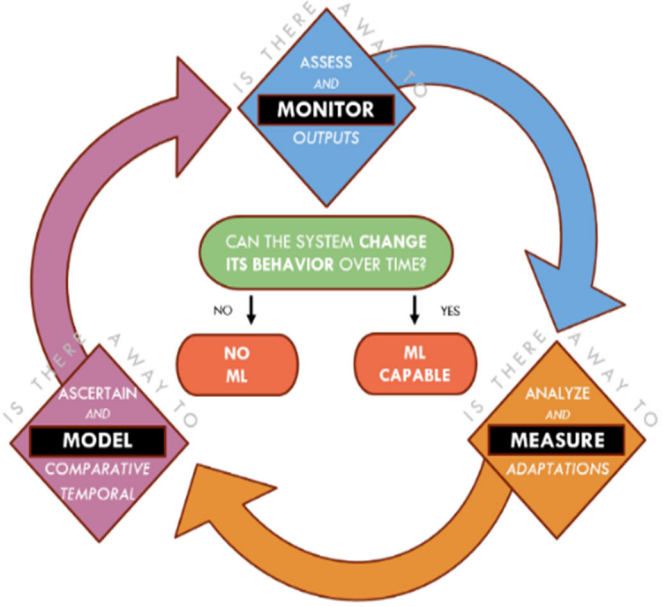

MONITOR: *Cameras, Microphones, Proximity Sensors, and Hard Drives for temporal or iterative changes*

MEASURE: *Characterize CPU/Bios Clock/System Policies, and Parameters changes in timing, performance, and other adaptations as they progress*

MODEL: *Correlate System Software (Algorithms/Apps), Logs (Temporal Logs) and Archives (Historical Data)*Throughout this process, the robot uses these sensors (*cameras, microphones, and proximity sensors*) to *monitor* activities. The robot then *measures* each output to validate they meet established parameters and then *models* them to inform optimization of performance. The robot uses these outputs to perform many functions. For example, it uses optical sensors to identify and recognize objects and ultrasonic sensors to monitor distance between objects. It orients these calculations relative to other objects in the room. Finally, it models the available options to discern which opportunity will provide the best pathway forward. The fishbone diagram below illustrates the available options and analyses.

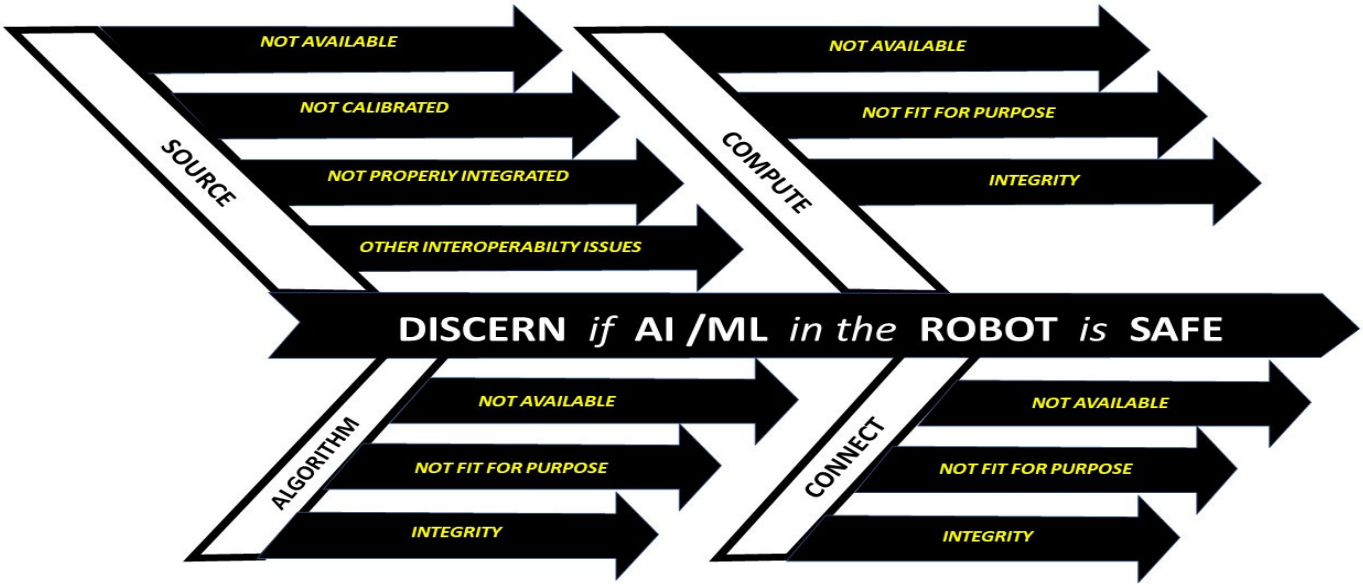

STEP TWO-ASSESS: For products where AI or ML have been found and that may have safety issues, subject matter experts (*SME*) assess what capabilities can be created by the components found in the screening process and establishes a test plan to characterize the AI/ML capabilities relative to their function within the system. Assessing AI/ML in accordance with its role within the product affords the opportunity to effectively test and evaluate these technologies.

During this step the function of each capability is determined and characterized by the features that it provides the product. In some cases, AI/ML is the product and in other circumstances it supports a product, by providing a feature that it relies upon. In other situations, AI/ML manages inherent risk within a product to ensure it operates within prescribed parameters. In all cases the AI/ML creates dependencies and, as such, the impact and resulting reliance on it must be evaluated in the analysis step.

STEP THREE - ANALYZE: To evaluate the system, each AI/ML capability's functionality must be analyzed for the contributions it makes to the product. During the analysis step, a determination is made whether AI/ML technologies contribute to a safety hazard. As previously outlined, the screening step has already identified AI/ML components and the assessment step has categorized current capabilities. In this third step, not only are the capabilities tested, and the system evaluated, but also the TE process examines the inputs and outputs in a variety of environments to discern if these technologies were contributing factors to creating the safety hazards.

The first stage of the analysis of the product is to document the problem. By establishing the requisite methodology, a scientific approach determines the contributing factors that cause the

safety incident or hazard. The following outlines this process with the suggested framework upon which to conduct a thorough scientific analysis:

- **Validate/verify known problem and specific hazards that create safety concerns in a controlled environment. This may include recreating conditions of a safety incident.**
- **Create test data to ensure consistent inputs throughout the TE process.**
- **Document configuration testing to establish a baseline to ensure testing consistency.**
- **Establish a process to ensure a consistent approach upon which to characterize how the capability acts and reacts throughout test iterations. Observe and record each influence and resulting effect.**
- **Repeat, as appropriate, to verify the test and validate the evaluation.**

The next stage of the analysis process is to characterize the functional capability at each stage of its lifecycle to determine the dependence on the design in isolation (*phase-I test capability*) and in conjunction with the other elements of the system (*phase-II evaluate integration*), and throughout the deployment of the product *(phase-III discern environmental adaptations*). This process will determine each function's contributing factor to the safety incident and isolate the AI/ML technologies' contributing factors to the safety concerns. The evaluation should include characterizing the influence of AI/ML design on the product, the implication/reliability of the product's AI/ML capabilities, and the characteristics and effects of product design in a variety of environmental conditions to include operating the product in reasonably foreseeable ways.

CONTRIBUTING FACTORS

- ***Identify if AI and/or ML was a contributing factor to the incident***
- ***Implications of what hazard each AI and/or ML capability contributed to***
- ***Impact of AI and/or ML technologies contribution to safety incident***

In our robot-use case example, the function of the sensors is to identify the way forward, while serving to elevate awareness. Sensory perception is just one example that provides the robot guidance essential for mobility. In this application, it is an AI capability that elevates situational awareness of obstacles that impede its path, as well as manages risk, by identifying situations that fall outside of safe parameters, like swimming pools, busy roads, and stairs that drop off abruptly. As the robot is informed by these sensors, it can determine which pathway provides the greatest propensity for success in order to avoid obstacles, and minimizes potential dangers that could impede progress.

Throughout the process, it is important to determine where the technology serves as both a function and feature within the system. The robot can use AI/ML to analyze each information stream (*e.g., environment, subcomponent performance)* to determine the factors contributing to potential unsafe conditions. For applications where AI/ML's role in the robot is to manage risk,

the AI/ML will discern whether and how component and system performance align with operating conditions. Ultimately, the objective of ongoing monitoring/measuring efforts by the robot identifies and characterizes the contributing factors that create hazardous conditions that cause safety mishaps, which are addressed in step four. Analysis of the Robot would include consideration of how the robot uses AI/ML in the examples above about potential unsafe conditions and risk management.

STEP FOUR – MONITOR/MEASURE: The adaptive nature of ML capabilities contributes to transformation of products that are in a constant state of evolution. To address this, parameters must be established in the assessment and analysis steps to provide benchmarks for identifying performance that is migrating to unsafe conditions. Through this step, AI/ML's impact is identified and evaluated for unforeseen and/or unanticipated outcomes, particularly those associated with migrating out of the parameters of safe operating ranges.

It is helpful to obtain documentation from manufacturers prior to evaluation to garner the historical context that captures parameters related to design intent and past performance, relative to present actions, which could impact the safety of a product. These efforts provide the origin and reference points to conclude when the AI/ML transforms beyond the parameters established for acceptable safe operations. By continually evaluating the product, we can determine when it will evolve beyond the current parameters and require a new assessment and analysis, given the current conditions and application of the product.

Ultimately the four-step process to screen, assess, analyze, monitor and measure is the core of the applied AI/ML TE program. The following three-phase approach looks at an incremental engagement that examines the impact and resulting implication of AI/ML in consumer products. In this progressive process, Phase-1 looks at only a specific AI/ML capability as it exists in the product; Phase-2 looks at all the AI/ML capabilities and how they are integrated into a system/product; and Phase-3 looks at the impact of the variety of environments and effects of varied data that are fed to the product. These progressive phases become increasingly more complex, costly, and time consuming, and they are becoming exponentially more difficult in both breadth and depth:

- **PHASE I (CAPABILITIES):** ***specific functions/features will be assessed/analyzed to discern if AI/ML is contributing to unreasonable hazards that are creating the conditions to cause a safety concern.***
  - Examine a specific AI and/or ML function or feature within a consumer product.

- **PHASE II (SYSTEM):** ***the system in which the AI/ML capabilities exist will be assessed/analyzed to determine whether there are contributing factors which create product safety concerns.***
  - Examine the AI/ML capabilities as integrated into the system of a product.

PHASE III (CONDITIONAL): ***the facts/data that are ingested by the system, as well as the interpretation of the outputs, and the environmental considerations are evaluated to determine if the data veracity, unforeseen use, or unintended consequences arise as a result of AI/ML technologies within the product.***

- Extensive evaluation of the product in a variety of applications to determine the potential of AI/ML to contribute to hazards in the context of reasonably foreseeable conditions.

RECOMMENDATION

Ultimately, a scientific approach is essential to identify if the basic ingredients necessary to create AI/ML capabilities exist, assess their implications, analyze the impact they have on the product, and monitor and measure AI/ML capabilities over time to determine if they contribute to product safety hazards for consumers. The proposed three-phase approach outlined in the "Way Forward" deconstructs the complexities of AI/ML through a TE program that looks at the contributing factors of each capability. Through this process, we can determine if AI/ML creates unreasonable hazards that contribute to known safety concerns when integrated into the consumer product.

This approach will require extensive collaboration, particularly given CPSC's limited resources. Fortunately, there are many organizations working to create voluntary standards like UL, ISO, ASTM, CTA, and IEEE. Continuing to work collaboratively with them and federal agencies like NIST who is working to refine the RMF is foundational to the TE program. Ongoing contributions to public-private partnerships like ACT-IAC will ensure the creation of a collective understanding of what AI is, as outlined in their AI Primer[16]; the best way to operationalize it, as illustrated in their AI Playbook[17]; and how to qualify and quantify the value as it has been exemplified in the Ethical Application of AI (EAAI) manual[18]. To this end, it is recommended that CPSC focus at least initially on the following collaborative steps:

- Contribute to ACT-IAC's AI/ML Use Case Repository
- Participate in NAII's talent development program
- Collaborate with the NAIIO to capture lessons learned and best practices.
- Continue to participate on the Federal AI Interagency Working Group
- Collaborate with and contribute to the NIST AI/ML Workshops

[16] Artificial Intelligence / Machine Learning Primer: *https://www.actiac.org/system/files/Artificial%20Intelligence%20Machine%20Learning%20Primer.pdf*
[17] AI Playbook for the U.S. Federal Government: *https://www.actiac.org/system/files/AI%20Playbook_1.pdf*
[18] Ethical Application of Artificial Intelligence Framework: *https://www.actiac.org/system/files/Ethical%20Application%20of%20AI%20Framework_0.pdf*

- Continue to participate in voluntary standards development, ensuring a focus on safety is maintained