Proceedings of the INMM Virtual Annual Meeting
July 24 – July 28, 2022

Inspecta 1.0: Development of an Architectural Roadmap

Heidi Smartt, Sydney Dorawa, David Hannasch, Philip Honnold, Nathan Shoman, Alexander Solodov, Katie Spence*
Sandia National Laboratories, Albuquerque, NM, United States
* EastSideWest Consulting LLC

ABSTRACT

Sandia National Laboratories is designing and developing an Artificial Intelligence (AI)-enabled smart digital assistant (SDA), Inspecta (International Nuclear Safeguards Personal Examination and Containment Tracking Assistant). The goal is to provide inspectors an in-field digital assistant that can perform tasks identified as tedious, challenging, or prone to human error. During 2021, we defined the requirements for Inspecta based on reviews of International Atomic Energy Agency (IAEA) publications and interviews with former IAEA inspectors. We then mapped the requirements to current commercial or open-source technical capabilities to provide a development path for an initial Inspecta prototype while highlighting potential research and development tasks. We selected a high-impact inspection task that could be performed by an early Inspecta prototype and are developing the initial architecture, including hardware platform. This paper describes the methodology for selecting an initial task scenario, the first set of Inspecta skills needed to assist with that task scenario and finally the design of Inspecta's architecture and platform.

INTRODUCTION

The IAEA Department of Safeguards is responsible for verifying international nuclear safeguards agreements. The mission of international safeguards is "to deter the spread of nuclear weapons by the early detection of the misuse of nuclear material or technology. This provides credible assurances that States are honouring their legal obligations that nuclear material is being used only for peaceful purposes" **Error! Reference source not found.**. The implementation of international safeguards is unique for different states, as they are based on sovereign agreements between a State and the IAEA, as well as from facility-to-facility as determined through a safeguards agreement's facility attachment. Safeguards activities at a nuclear facility are also based on state factors and the IAEA's technical objectives as defined in the Annual Implementation Plan.

Despite these variations, there are many common and repetitive inspection activities performed by inspectors such as reviewing facility bookkeeping, physically inspecting and maintaining safeguards equipment, taking measurements and samples, examining and verifying seals, item counting, reviewing surveillance images, verifying design information, and generally observing a site for discrepancies. These inspection activities are often mentally and physically challenging and thus may be susceptible to human error. Additionally, there is an upward trend in the number of responsibilities for international safeguards inspectors. This increase in responsibilities is a direct result of 1) an

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. SAND2022- C

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

increase in the quantity and types of nuclear facilities under safeguards related to the development of novel nuclear fuel cycles; 2) an increase in the global number of significant quantities of special nuclear materials due to the longevity of safeguards for waste products and spent fuel; and 3) a push for inspectors to move from a traditional role of "auditors" in the field to more investigative roles in which activities are defined via technical objectives in pre-defined safeguards criteria. Even with these increased responsibilities, inspectors still have limited time in facilities and must work as efficiently and effectively as possible. Efforts to enforce obligations associated with the Additional Protocol (AP) will add to this workload.

Artificial Intelligence and its underlying algorithms are prominent and increasingly present in our everyday lives, i.e., cars with automated driver-assistance, online vendors suggesting future purchases, voice-assisted smart home controls, AI/robotic vacuum cleaners, and SDAs like Amazon's Alexa[1]. Integrating these advanced capabilities with international nuclear safeguards inspection processes could increase the effectiveness and efficiency of safeguards activities, especially for those tasks that are tedious, challenging, and prone to human error.

We are developing a prototype for an AI-enabled SDA for safeguards inspectors to support their increasingly challenging task requirements, named Inspecta (abbreviated for "International Nuclear Safeguards Personal Examination and Containment Tracking Assistant"). Inspecta is similar in function to Alexa or Siri[2] as it can aid with tasks such as note-taking, alarms, and timers, but will also have safeguards-specific task capabilities like using optical character recognition (OCR) to read seal numbers. Inspecta will reside on a small, portable, and potentially wearable device that will primarily interact with an inspector verbally, with some capabilities to display information on a screen as appropriate. Note that Inspecta is intended to work alongside humans and not as a replacement for inspectors.

In this paper, we will share the methodology used to determine a high-impact inspection task that would benefit from Inspecta's assistance, describe the technical capabilities ("skills") required to perform this inspection task, and describe the initial architecture of the Inspecta 1.0 prototype.

METHODOLOGY TO DOWN-SELECT TASK

In 2021, we started with precisely defining what Inspecta would be, developing a list of high-impact tasks performed by IAEA inspectors, and creating a list of technical skills an SDA would need to support an inspector in their tasks. We drew on three sources of information as part of this process: (1) comprehensive lists of IAEA safeguards inspection tasks, (2) IAEA publications to identify challenges [2, 3, 4, 5], and (3) interviews with 8 former IAEA inspectors and subject matter experts with a focus on identifying which inspection tasks were most challenging, tedious, or prone to human error. We extracted tasks that were identified the greatest number of times during these interviews as

[1] https://developer.amazon.com/en-US/alexa
[2] https://www.apple.com/siri/

(1) surveillance review, (2) transcription, (3) information integration, (4) seals examination and verification, and (5) spent fuel verification. Transcription and integrating information are often part of the other identified tasks and thus we focused on surveillance review, seals, and spent fuel verification. Other R&D programs utilize machine learning for surveillance review and using floating robots for spent fuel verification [6]. Therefore, the task chosen for Inspecta 1.0 was examining[3] CAPS seals (hereafter called "seal-examination"). The skills developed for this task can be re-applied for other tasks in future iterations of Inspecta.

TECHNICAL CAPABILITIES/SKILLS

The seal-examination task is important but tedious for IAEA inspectors. Tamper-evident metal CAP seals (Figure 1) with a numeric identifier are attached to containers after the contents have been verified. An inspector will be escorted by facility personnel to the material holding location, find seals to examine, compare the seal number on the item with the seal number on a paper list, and mark that the seal has been examined and confirmed. The inspector also physically inspects the seal and seal wire for signs of tampering, pulling on the wire and seal to ensure proper connection to the container. A small set of seals may be selected for removal and verification at the IAEA headquarters; this selection process is performed using a statistical algorithm that informs how many and which seals should be removed and replaced.

Figure 1: CAPS metal seal, image: IAEA, 2020.

Inspecta can provide varying levels of assistance to the seal-examination task to support inspectors. At one extreme, Inspecta could fully automate the seal-examination task alongside an IAEA inspector. This full automation could use robotics and indoor navigation to locate the area with seals, use a robotic manipulator arm and object detection to find and grasp a seal, apply OCR to acquire the seal number and compare to a local database, physically confirm attachment by tugging on the seal and wire, employ anomaly detection to identify signs of tamper, and communicate with the inspector

[3] The CAPS seal cannot be fully verified in the field and is removed and taken to IAEA headquarters for full verification. We'll use the term "examination" to include visual and physical inspection and comparing a seal number to a seal inventory list.

using speech synthesis, speech recognition, and information recall. Figure 2 illustrates skills that would be applicable in this fully automated case.

For Inspecta 1.0, the project scope is significantly narrower than for the fully automated case. Focus is on skills that would directly aid an inspector while still requiring their full involvement in the seal-examination task. Namely, OCR is targeted to semi-automate the review of individual seals, and speech synthesis and speech recognition skills are targeted to simplify notetaking. With Inspecta 1.0, an inspector will still perform physical inspection. Other skills may be added to Inspecta in future versions.

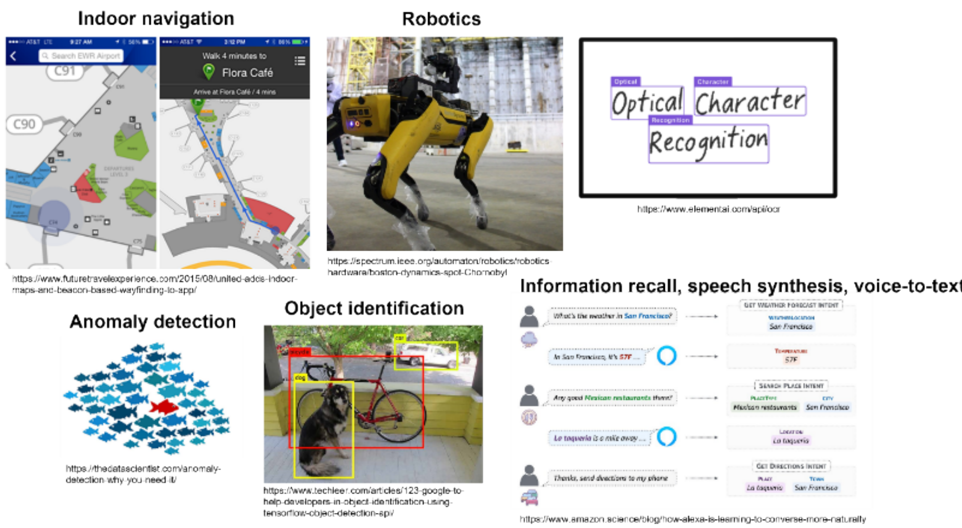

Figure 2: Skills applicable to "seal-examination" task.

INSPECTA 1.0 ARCHITECTURE

Before development of an Inspecta architecture, we considered high-level requirements that would drive the design. These were mainly centered on security, privacy, usability, and the ability to use and modify previously developed algorithms, libraries, and software components. Hardware that already included many of the input/output capabilities and device sensors was also considered. An early decision was whether Alexa, Siri, or the open-source Mycroft[4] could be used as a baseline to build from, but project constraints eliminated those platforms as options. Alexa and Siri are cloud-based and require connection to servers to work properly. Wireless connections in nuclear facilities are often not available (or reliable), either due to security concerns from the facility or due to signal dead zones throughout the facility. Further, both platforms are proprietary which would limit access to source code and could also be a security concern for both the facility operator and the IAEA. Privacy would be questionable since data is sent/received through cloud-based servers. While Mycroft is open source, there are few turnkey skills available, and relatively limited documentation available. For

[4] https://mycroft.ai/about-mycroft/

these reasons, construction of the SDA platform from the ground up was deemed the best development pathway.

Development of Inspecta 1.0 requires a development platform, a hardware device, and both a software application and machine learning algorithms loaded on a hardware device. As mentioned previously, we sought a hardware device that had integral input/output capabilities – namely, speakers, microphone, camera, and display, and chose a smart phone as a candidate. While cellular capabilities are often prohibited in nuclear facilities, there are existing seal readers that use modified cellphones where the antenna has been grounded to prevent communication. Smart phones also have various internal sensors for tracking and navigation, a feature that may be beneficial for future Inspecta versions. Smart phones are small, portable, and can be wearable, which is important since inspectors will be carrying the device throughout the day. Finally, smart phones have reasonable computing power and extended-life batteries and battery extension packs.

Current efforts are aimed at deploying Inspecta on an Android phone for a demonstration of the seal-examination task at the end of 2022. The Android platform was selected due to the ease of development and ubiquitous deployment of the operating system; there are roughly 2.3 billion Android devices that range from phones to embedded devices. This ensures that there are many hardware deployments options and extensive supporting documentation.

While the near-term development objective targets an Android phone, it is anticipated that actual devices allowed into a facility may vary from site to site. As such, Inspecta is being developed in Xamarin, which is an open-source application development platform based on .NET and C#. Xamarin allows the development team to write Inspecta in a shared code base that can be deployed simultaneously as an Android, iOS, and Windows application. This enables the team to write a single set of code that can be run on multiple different platforms without additional development effort.

Another key requirement for Inspecta is the prioritization of on-device learning. A critical constraint of machine learning use is the training of an algorithm. There are many different cloud providers that offer machine learning services; however, it is likely that data must remain on-device due to safeguards agreements. The Open Neural Network Exchange (ONNX) application program interface (API) is currently being used to develop models that can be used by the Xamarin-based Inspecta code. Specifically, models built in PyTorch and Tensorflow are converted to an ONNX format, which can run on the various deployment targets (i.e., Android, iOS, and Windows). Additionally, models are quantized to improve performance for embedded and mobile devices. Quantization [7] is the process of lowering precision of the machine learning algorithms' mathematical operations to improve performance at inference time.

Inspecta is being developed with a modular framework with reusable components that can be used for future capabilities, as illustrated in Figure 3. Currently, prototype modules for the speech recognition and speech synthesis have been created. The speech recognition module utilizes a modification of a state-of-the-art transformer architecture called Wav2Vec [8] whereas the speech synthesis uses native device capabilities that can be used offline. Tacotron [9] is currently being studied as a potential speech synthesis module. Additional efforts are focused on using Levenshtein distance [10] to match voice commands to available Inspecta actions.

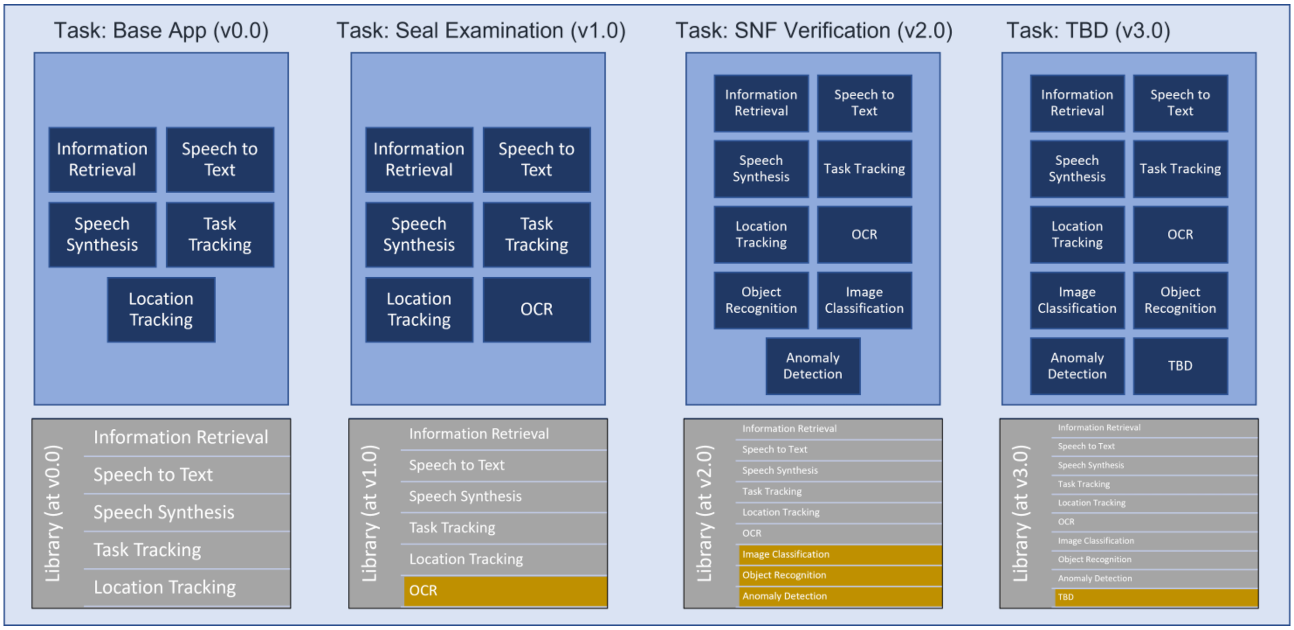

Figure 3: Modular architecture enables Inspecta to adapt to future needs.

SUMMARY AND NEXT STEPS

An AI-enabled SDA can be integrated into the process of international nuclear safeguards inspections to assist with mentally and physically challenging tasks and those prone to human error to increase the effectiveness and efficiency of inspections. In this work, we have identified and down-selected safeguards tasks that are mentally/physically challenging and prone to error, based on subject matter expert interviews and mapped these tasks to Inspecta skills and required technical capabilities to perform these skills. We have selected one task (seal-examination) to demonstrate feasibility in an Inspecta 1.0 prototype and will use the skills needed to perform that task as the basis for building up future functionality. We are currently developing the architecture, software application, and machine learning algorithms (selecting and modifying for the specific application) for speech synthesis, speech recognition, and OCR. We are working closely with a related NA-22 project "Trust in Auditory User Interface" which will help guide specific user interface design considering factors like privacy considerations associated with communicating sensitive information with Inspecta, inspector trust in the AI's decision-making process, and identifying features best suited for auditory, visual, or audio-visual interfaces.

ACKNOWLEDGEMENTS

The authors would like to acknowledge and thank the U.S. National Nuclear Security Administration (NNSA) Office of Defense Nuclear Nonproliferation R&D Safeguards portfolio for funding and supporting this research.

REFERENCES

[1] International Atomic Energy Agency. https://www.iaea.org/topics/basics-of-iaea-safeguards. Accessed 3/17/2021

[2] IAEA Safeguards, "Emerging Technologies Workshop: Trends and Implications for Safeguards Workshop Report." February 2017, Available at: https://www.iaea.org/sites/default/files/18/09/emerging-technologies-130217.pdf

[3] IAEA Safeguards, "Emerging Technologies Workshop: Insights and Actionable Ideas for Key Safeguards Challenges Workshop Report." STR-397, January 2020. Available at: https://www.iaea.org/sites/default/files/20/06/emerging-tehnologies-workshop-290120.pdf

[4] IAEA Safeguards, "Research and Development Plan: Enhancing Capabilities for Nuclear Verification." STR-385, January 2018. Available at: https://www.iaea.org/sites/default/files/20/06/rd-plan-2018.pdf

[5] IAEA Safeguards, "Development and Implementation Support Programme for Nuclear Verification 2020-2021." STR-393, January 2020. Available at: https://www.iaea.org/sites/default/files/20/06/dis-plan-2020-2021.pdf

[6] Adem Mutluer, Robotics Challenge Winning Design Helps Speed up Spent Fuel Verification, March 2019, https://www.iaea.org/newscenter/news/robotics-challengewinning-design-helps-speed-up-spent-fuel-verification Accessed 8/3/2021.

[7] Gholami, A., Kim, S., Dong, Z., Yao, Z., Mahoney, M. W., & Keutzer, K. (2021). A survey of quantization methods for efficient neural network inference. *arXiv preprint arXiv:2103.13630*.

[8] Schneider, S., Baevski, A., Collobert, R., & Auli, M. (2019). wav2vec: Unsupervised pre-training for speech recognition. arXiv preprint arXiv:1904.05862.

[9] Shen, J., Pang, R., Weiss, R. J., Schuster, M., Jaitly, N., Yang, Z., ... & Wu, Y. (2018, April). Natural tts synthesis by conditioning wavenet on mel spectrogram predictions. In *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)* (pp. 4779-4783). IEEE.

[10] Levenshtein, Vladimir I. (February 1966). "Binary codes capable of correcting deletions, insertions, and reversals". Soviet Physics Doklady. 10 (8): 707–710