



PRINCIPLES AND PRACTICES FOR THE RESPONSIBLE APPLICATION OF ARTIFICIAL INTELLIGENCE AT MOTOROLA SOLUTIONS

WHITE PAPER





Artificial Intelligence (AI) will improve the efficiency, effectiveness and safety of the Motorola Solutions user community. However, as a powerful, multi-faceted emerging technology, AI can have far-reaching unintended consequences if applied inappropriately, or executed with insufficient rigor and discipline.

This paper summarizes Motorola Solutions' policies and practices for responsibly applying AI in its public safety products and applications. Our goal is to ensure unbiased, fair, understandable, secure and reliable operation. We apply a human-centered design philosophy that is informed by copious customer research to guide our application of AI based on the foundational tenets of "human in the loop" for all consequential actions, and an emphasis on the use of human-centered design in the application of AI for focused, purpose-built solutions.

HOW ARTIFICIAL INTELLIGENCE BENEFITS OUR PUBLIC SAFETY USERS

To a layperson, AI is generally presumed to be computer technology that can broadly mimic a human in performing complex tasks. In practice, today's technology is not that advanced. Currently, AI can only emulate aspects of human intelligence, focused on narrowly defined tasks such as the ability to recognize faces, or perform a translation of audible speech to text. In the context of responsibly applying AI to the public safety workflow, these specific task applications are, in fact, more beneficial.

Motorola Solutions has been creating mission critical communications hardware and software for nearly our entire 91-year history. Through our acquisitions, strategic venture investments, and evolving research and development profile, we are assembling a comprehensive operating platform for public safety. This platform is the foundation upon which we can bring AI-driven outcomes to our customers. As the world creates exponentially increasing quantities of structured and unstructured data that characterizes conditions and events, we are applying AI to interpret this data, improving the efficiency, effectiveness and safety of first responders and command center personnel around focused tasks.

Our objective is to support humans throughout the public safety incident workflow by augmenting their decision-making with appropriate AI assistance. Some ways that we are applying AI in this regard include:

- Using AI to transcribe, translate, interpret and summarize speech and text during human interactions with a public safety professional or emergency call center.
- Utilizing AI to power natural language interactions (voice bots), enabling first responders to access critical information while allowing them to remain aware and attentive ("eyes up and hands free") at all times. Voice applications could include running a license plate check with a voice query/response; populating a form verbally rather than manually; or a periodic autonomous well-being check by asking a first responder to verbally acknowledge their condition. A key aspect of successfully executing voice interaction - informed by extensive experience building mission-critical voice solutions - is building the vocabulary of our user groups, e.g. 10-codes and the NATO alphabet, into the experience.
- Leveraging AI to "watch" the exponentially increasing video sources available to public safety, identifying unusual occurrences in real time and alerting a human analyst. These use cases could include noting the appearance of smoke, the formation of a crowd, a person crossing a security boundary, or identifying an individual matching a description such as a missing person or AMBER Alert subject of interest. AI also excels at searching historical video for items or persons of interest.
- Applying AI in the form of biometric identification (such as facial recognition) in public safety applications to identify individuals without identification or who are incapable of identifying themselves.

In all cases, these are tasks that humans perform manually today. We are simply using AI to automate the labor-intensive aspects for each.



THE NEED FOR RESPONSIBLE AI

Artificial Intelligence is new, complex, powerful technology. Applied carelessly, it can generate surprising and unintended results. Applied improperly or maliciously, the technological scaling of AI can amplify or institutionalize undesirable outcomes - a risk of particular import for public safety, given its societal impact.

Responsible AI simply means that when applying AI in products and services, one must consider the ramifications (good, bad, or unexpected) and constrain the outcomes to ensure proper results. Fundamentally, AI powered solutions must be unbiased, secure and trustworthy. Motorola Solutions approaches responsible AI based on our guiding policies and principles, supplanted by the underpinning human-centered design, development, validation, and performance monitoring processes that ensure our AI solutions adhere to the policies.

POLICY AND PRINCIPLES FOR APPLYING AI RESPONSIBLY

Over decades of designing mission critical systems, Motorola Solutions has developed substantial design practices expertise founded in user research, iterative user-inclusive design validation, and domain-specific human factors engineering. We annually accumulate thousands of hours of in-field and in-situ observational user research (ethnography). We apply this intimate understanding of user needs to a human-centered design process that results in purpose-built, focused solution designs specifically optimized for the tasks at hand. We then develop prototypes and verification frameworks that we test and validate iteratively with users to perfect the design and, ultimately, its implementation. Finally, we apply specific human factors disciplines to our products to ensure that the technology is well-tuned to the needs of first responders operating across their daily workflows.

We are applying the same focused research and design-led approach in the application of 'purpose built' AI solutions to ensure that we are aligning the capabilities with specific human needs (no more, no less).

Our policy for applying AI starts with three foundational tenets that are fundamentally focused on our users and fields of use:

- Human in the Loop
- Focused application resulting in purpose-built solutions
- Applying mature AI

HUMAN IN THE LOOP

Fundamentally, our approach is to augment human decision making while never displacing or disintermediating human judgement. In this sense, our systems are advisory and will never take AI-generated consequential actions on their own. In other words, there is always a "human in the loop" to make the final determinations on substantial decisions. Studies indicate that the best results are achieved with a combination of AI and human experts¹.

Our user-centered research methodologies have led us to a specific set of human factors disciplines that we refer to as High Velocity Human Factors (HVHF). This recognizes that the more stress and duress an individual is experiencing, the less cognitive capacity they have to apply to anything other than the specific event.

For example, if a police officer has deployed his or her weapon, there is little else they can (or should be) focused on other than the threat at hand. The paradox of HVHF is that the more a user could benefit from technology - including AI recommendations, for example - the less mental capacity they have available to leverage it.

We are investigating the selective application of AI to detect these situations and understand the context so that we can automatically adapt the operation of the technology on behalf of the user. For example, detecting the condition via biometric data, voice inflection, and/or background audio event detection that leads to activation of streaming body worn video, prioritization of communications, alerting the command center and nearby officers, or similar.



FOCUSED APPLICATION AND MATURE AI

Our two remaining foundational tenets - focused application, and mature AI - are connected.

Many respected and capable companies across a variety of industries have applied substantial thought and effort into the considerations of how to best apply AI in a general fashion, across a diverse set of applications, users, and environments^{2,3,4}. Applying AI in a general manner is extremely difficult, due to the wide variety of possible circumstances and user personas.

Given our emphasis on public safety, Motorola Solutions can focus this particular AI challenge. We are creating "purpose-built" AI solutions focused on specific tasks within the public safety workflow. Our immersive human-centered design practices result in a thorough understanding of the end-to-end process, and awareness of bottlenecks within the system where AI can accelerate or improve the process.

We are not creating general AI utilities, tools and frameworks for others to leverage in building their solution. Rather, we are applying proven AI technology in a focused and collaboratively developed way. We are placing it in the hands of a known population (e.g., first responders, or dispatchers). We are augmenting their decision-making, not displacing them. Finally, we are applying the technology to a narrowly defined and well-understood use case.

Critically, the narrow scope of the application allows us to choose the simplest, most specific, and most mature underlying AI technologies. We can manage and curate training data for the specific applications and environments implied by the solution, and to constrain the scenarios and use cases across which the solutions must be validated.

Motorola cannot ensure or enforce our customer's compliance with legal and ethical standards in the application of AI, but our products provide access controls based upon user authentication credentials and access permissions to allow our customers to enforce user compliance to operational policies. Our products will maintain an audit trail of user operations and where possible, constrain the use of AI powered capabilities within the context of a workflow. For example, submitting an image of a face (probe) for search against a gallery may only be accomplished in the context of an active investigation by authorized individuals. When a 'match' is surfaced, this must be verified and signed off by one (and sometimes two, depending upon local policy) human experts. Audit trails will not only keep track of the specific operations and user identity but will also retain any data used in the operation (e.g., introduction of a probe-image for search).



FUNDAMENTALS OF RESPONSIBLE AI

Underlying our foundational tenets designed for the unique needs of public safety, there are generally acknowledged universal considerations when applying AI that Motorola Solutions addresses:

- Bias and fairness
- Understandability and transparency
- Privacy
- Reliability and security

BIAS AND FAIRNESS

Bias, in general, is a prejudice for or against one thing, person, or group compared with another, usually in a way considered to be unfair. For humans, bias can result from a number of factors, including a person's environment, training, condition, or experiences. In AI, bias occurs when the output of the AI process results in inconsistent treatment across a group. For example, being able to identify faces more accurately for one demographic (characterized by race, gender, age, and physiology) than another, or making decisions about how a person should be treated (such as a hiring decision) inconsistently across demographics.

Much like humans learn iteratively through experience, AI algorithms learn by iteratively 'training' with representative data, developing the ability to identify discriminating characteristics across this data set. For example, in a facial recognition application, the AI system is exposed to a series of pictures of different and identical faces that are tagged so the AI can know whether given images are of the same individual or not. In this way, the AI effectively learns what characteristics distinguish different people. (This is very similar to how humans develop their ability to recognize and differentiate patterns in the senses of sight, sound, smell, taste, and touch.) Voice interaction systems are essentially trained in the same manner using tagged verbal data. The completeness and accuracy of the AI system's 'learning' is a function of the quality, diversity, correctness, and volume of the training dataset.

In general, data needs to be completely representative of the context that the AI system

will operate within (population, vocabulary, etc.) and there must be sufficient data available for the AI to develop the ability to fully discern uniquely identifying characteristics.

Bias can be present within training data in two ways: either the data is unrepresentative of the population that the AI will operate within, or it reflects past or existing prejudices. For instance, the first case might occur if an AI system trains on a set of voice samples that has more females than males. The resulting natural language interaction system would be worse at recognizing male speech. An example of the second case might be if a recruiting tool dismisses a large proportion of female candidates because it was trained on historical hiring decisions, which were dominated by males over females.

AI is fundamentally amoral as it is not influenced by human discriminatory tendencies, emotions, distractions or fatigue. Thus, AI has the ability to be much more deterministic, and less biased than human counterparts. Furthermore, as the AI system is used and exposed to more data in its population set, its accuracy will continually improve.

How Motorola Solutions Supports Fairness in our AI

Motorola Solutions thoroughly evaluates the data employed in training our AI algorithms to ensure that sufficient quantity, quality and diversity exists across the dataset to properly train the algorithm for its intended purpose and operating environment.

We thoroughly validate the operation of the trained algorithms with a representative and diverse set of test cases that are applied across a range of operational conditions. We also test and retest our products in actual customer environments.

In fielded operation, our systems generate telemetry that we can continuously monitor to identify performance issues, as well as any inconsistent or undesirable behaviors. For example, as we discover optimizations and refinements that improve accuracy we can retrain the model with an enhanced data set (e.g., a new set of slang terminology or speech accent that was not initially known/anticipated) and deploy that universally across our entire user base.

UNDERSTANDABILITY AND TRANSPARENCY

Understandability simply means that the behavior and outputs of an AI system must be readily explainable by those who provide it. This 'why' component is an essential characteristic of a system in order for users to interpret and trust the outputs produced, and for society to trust the tools that first responders are applying on their behalf.

In traditional systems, explaining behavior is generally a straightforward exercise of reading the software programs to understand what the programmer instructed the system to do. Trained AI systems do not fundamentally have programmed or codified behaviors. Their operations are a function of the training dataset used. Variations in output may come from the system's training or may be the manifestation of traditional software bugs in the underlying implementation of the model.

How Motorola Solutions Supports Understandability in our AI

- Motorola Solutions maximizes our ability to explain the operation of our systems by adopting mature, testable AI components that are as simple as possible for the task at hand. For example, by knowing the verbal commands needed for an interaction that the system will support, we can greatly constrain the verbal phrases and words (intents) that need to be understood.
- By leveraging mature, well characterized and understood implementations and model frameworks with proven reliability, we increase the likelihood that the simplest explanation is most likely to be correct. Simple systems are much easier to test, understand and explain than complex ones, and also boost the overall reliability of the system. Often, these components will come from services in widely deployed cloud service platforms operating at scale.
- We ensure that our systems generate operational performance data that we can monitor and review on a regular basis to assess efficacy of operation.

PRIVACY

In the context of AI applications, the issue of privacy primarily involves securing and managing the data associated with the system, including training data (input to the system) and any outputs that the system produces in operation.

Training data may contain Personally Identifiable Information (PII), which has to be managed in accordance with appropriate information security policies. As established in the bias and fairness discussion, it is essential to work with representative data when training AI. Consequently, Motorola Solutions will almost certainly be leveraging samples of live data (containing PII) in some instances.

How Motorola Solutions Supports Privacy in our AI

- Where possible, Motorola Solutions will work with anonymized data (no assignable PII content) at the source, synthesized data (by machine methods or through controlled customer interactions such as training exercises), or accumulate our training data from publicly available sources.
- We utilize tools and frameworks that facilitate privacy-sensitive training by encoding general patterns rather than facts about specific training examples wherever possible. Techniques such as differential privacy and federated learning offer strong mathematical guarantees that models do not learn or remember the details about any specific user. Furthermore, no training data is distributed with our products.
- In the USA, Motorola Solutions employs Criminal Justice Information System (CJIS) compliant data facilities, approved personnel, and stringent practices to manage the data that we house for training and testing algorithm purposes. Outside of the USA, where it is necessary to apply localized data for this purpose, we will adhere to all country specific regulations and practices that might entail establishment of local presence and a secure data repository that respects considerations such as data sovereignty.
- When securing usage or deployment-generated operational outputs of AI enabled solutions, we rely upon the same

stringent cybersecurity practices that we apply to all of our products⁵. We build in all of the necessary security controls, auditing, and practices necessary to enable our users to secure and manage sensitive data (e.g., Law Enforcement Records) and this same fabric applies to outputs generated by AI. Where possible we leverage industry best practices and scale (e.g., government cloud based platforms and environments).

RELIABILITY AND SECURITY

As with any item that public safety users leverage for day-to-day operation, AI based solutions must be as reliable as possible. Motorola Solutions has a history of delivering mission critical solutions and a culture of ensuring quality, performance, and integrity in our products. This is accomplished by following disciplined development, validation, deployment and life-cycle management processes and employing comprehensive product testing practices.

When introducing changes to solutions in the public safety space, we employ an incremental discipline. Starting with a thorough understanding of users and their environment, we progress through a collaborative creation process where we iteratively develop, test and refine capabilities incrementally with continual customer feedback.

How Motorola Solutions Supports Reliability and Security in our AI

- We ensure that new capabilities don't eliminate, impair or alter existing functions so that in the moment of need nothing previously relied upon or learned as 'muscle memory' is compromised.
- We introduce new capabilities in controlled and tightly monitored circumstances (e.g., training facilities and environments for models, select test users, etc.) and employ focus group feedback. In selecting these target environments, we consider the nature of the AI capability being applied and ensure that we choose diverse environments to maximize the testing surface variation.
- We employ standard incremental rollout methods across our user base such as Beta testing.

- We leverage continuous integration / continuous deployment development practices that allows us to deploy changes, fixes, and new capabilities quickly and incrementally. This is especially important in the context of retraining AI functionality.
- We build in system monitoring and telemetry that we can threshold or monitor regularly to ensure that system is operating as intended.
- Our products build in compliance mechanisms to apply role-based access controls and create audits, allowing our customers to enforce compliance and accountability across their users.

SUMMARY

Motorola Solutions is responsibly and incrementally employing AI to assist and augment our users to help them be more efficient, effective and safe. We are doing this by leveraging proven mission critical research and design principles guided by the fundamental tenets of human in the loop for consequential decisions and focused solutions that leverage mature AI.

We are developing AI aimed at customer outcomes that are familiar and consistent with outcomes we've previously enabled in a more manual fashion with other technologies. AI simply enables these outcomes in a way that is more efficient, more effective, and safer. We can easily measure the effectiveness and accuracy of AI-driven solutions relative to more traditional methods of achieving the same outcomes. In this way, our AI solutions are anchored in and measured against widely accepted, culturally and ethically appropriate methods that support fairness, understandability, privacy, and security.

THE UNDERSTANDABILITY POTENTIAL OF MACHINE LEARNING

Where possible and applicable, we apply solutions that provide more interpretability for our customers. There are exciting advances being made rapidly in the arena of AI interpretability via machine learning.

Fundamentally, traditional software uses human readable rules to map inputs to outputs. It is possible to use machine learning to learn these rules directly or to decipher the implicit logic used by the trained AI model to map inputs to outputs. We ensure that our machine learning models are deterministic and where possible interpretable as rules operating on the input.

When machine learning is applied, we will deploy trained and tested implementations rather than field-based machine learning wherever possible. When we do apply machine learning it will be for low-level cases such as allowing the solution to adapt to its environment (e.g., using machine learning on a video camera to estimate scene geometry and constant conditions). This allows us to curate the training data set to the specific objective at hand and validate the operation of the system through in-house testing prior to deployment.

REFERENCES

1. Face recognition accuracy of forensic examiners, superrecognizers, and face recognition algorithms; Philips, Yates, Hu, Hahn, Noyes, Jackson, Cavazos, Jeckeln, Ranjan, Sankaranarayanan, Chen, Castillo, Chellappa, White, O'Toole; Proceedings of the National Academy of Sciences; June 2018
2. The Future Computed – Artificial Intelligence and its Role in Society; Microsoft Corporation; 2018
3. Responsible AI Practices; Google
4. Everyday Ethics for Artificial Intelligence; IBM Corporation; September 2018
5. Motorola Solutions Products, Solutions, and Services Foundational Cybersecurity Policies; Motorola Solutions; Version 1.0; October 2017

To learn more about empowering public safety with AI, visit:
motorolasolutions.com/viqj

