Olive
99 E Main Street
Columbus, OH 43215

September 15, 2021

Re: Olive AI, Inc.'s Response to the National Institute of Standards and Technology Request for Information – AI Risk Management Framework

Olive AI, Inc. ("Olive") submits the below in response to NIST's Request for Information regarding the AI Risk Management Framework (NIST-2021-0004).

Olive's AI workforce is built to address healthcare systems' most burdensome issues -- delivering hospitals and health systems improved efficiency, increased revenue, lower costs, greater productivity, reduced errors and increased capacity. Patients feel lost in the system today and healthcare employees work in the dark due to outdated technology that creates a lack of shared knowledge and siloed data. Olive is designed to drive connections, shining a new light on outdated healthcare processes that stand between providers and patient care. Olive uses AI to reveal life-changing insights that make healthcare more efficient, affordable and effective. Olive is improving healthcare operations today, so everyone can benefit from a healthier industry tomorrow.

In an industry that continues to rely on outdated technology like fax machines, the power and potential of AI is not more evident than it is in healthcare. From lowering costs to improving patient care, AI has proven to not only improve outcomes for all industry participants, but also save lives. An AI Risk Management Framework as proposed by NIST provides innovative developers with the guidance needed to best assess how to appropriately manage risks, allowing the full potential and opportunity of AI to be unlocked.

Olive welcomes the development of an AI Risk Management Framework to establish guidelines and standards for AI risk management in order to promote and create consistency and confidence in the development of AI across various industries. Olive has summarized its feedback to NIST's Request for Information into the following points:

1. **Specific Industry Considerations**

AI is unique in that many of its proposed and potential uses have not yet been conceived; however, that does not eliminate the need for the proposed AI Risk Management Framework. Instead, it stresses the need for a flexible and industry agnostic framework that allows companies to address the unique needs and risks (current & future) of not only their industry but also the

proposed use of AI. For example, while there are general AI risks within the healthcare industry, even those risks widely vary depending on application (i.e. billing processes versus patient care). Any AI Risk Management Framework should provide guidelines on how to weigh these various considerations related to risk.

2. **Pre-Trained AI Models & APIs**

While many conversations around AI currently focus on the underlying development of AI tools, any AI Risk Management Framework should also take into consideration and provide guidance to organizations on the management of risks related to the implementation of pre-built AI tools (e.g. pre-trained, large parameter deep learning models). Unforeseen issues arise when users of these pre-built AI tools do not have the ability to validate the ethical treatment of risks and adherence to a "do no harm" stance. Thus, users need a way for the Framework to guide organizations on how best to adopt pre-built AI tools in a manner that minimizes risk while enabling fast-paced development. That is, the AI Risk Management Framework should arm organizations with tools to make good use of AI tooling, without becoming a barrier to entry or impacting economic viability.

3. **Managing & Addressing Bias**

One of the most concerning AI risks is bias that results in unplanned and unexpected negative outcomes. The largest drivers of biases in AI are the inherent biases in the training data made available during the creation of the AI itself. Recent and numerous examples exist where well-meaning teams unintentionally built heavily biased AI systems, in large part because of unknown bias (due to the lack of quality and quantity) in the training data. Without tools to identify biases in training data, such problems are only caught after considerable expense, which only serves to encourage well-intentioned organizations to obfuscate the degree to which their product has implicit bias, if they are even aware of it. Given that robust AI based solutions require sufficient and quality data upon which to be built, consideration should be given to data protection regulations and frameworks to ensure they promote and encourage the exchange of data to meet this need. The NIST AI Risk Management Framework should offer guidance and factors for organizations to consider to eliminate, mitigate, and manage bias in the development and implementation of AI tools, including where sensitive data points (such as race, gender, etc.) may be used by AI tools for beneficial outcomes.

4. AI Development Lifecycle

In developing an AI Risk Management Framework, consideration should be given to creating a baseline and defining the AI Development Lifecycle. Similar to industry standards around the software development lifecycle ("SDLC") or data protection impact assessments, this would ensure that the proper risks are being considered at each stage of development and in a timely fashion.

For example, there are often unique or higher risk tolerances for AI in development, as opposed to what would be tolerated in a production environment. This should be a consideration in the development and flexibility of an AI Risk Management Framework. Additionally, similar to how the management of risk adapts to the expected usage of a software product, so too should the AI Risk Management Framework guidelines evolve with drift potential and factors that may impact accuracy or reliability.

Model and metric design are essential to a project's success and integrity. There currently is no single framework or blueprint for project/model architecture, and the art of ML/AI development makes such a unified offering unlikely. For example, the design complexity is increased significantly when working with time-aware modeling for forecasting projects when compared with binary classification of tabular data. In addition to this, decisions about metric design and calculation are equally impactful. Consider the case of how using accuracy as the reporting metric for classification models with imbalanced class sizes can lead to a false sense of success. In such cases, the use of alternative metrics that are better suited for imbalanced classes provide better context for the model performance. The culmination of these design decisions can result in a model/algorithm that performs well within the test constructs but fails when implemented in a live setting. If monitoring is set to test initial design metrics, tests might continue to appear performant while the actual performance gaps are missed. Time, effort, and resource expectations need to include understanding of the complexity of design and testing required to implement a successful solution. As an AI tool remains in production, there then should be a periodic review cycle to ensure it continues to perform as expected.

Guidance provided by the Risk Management Framework on these issues throughout the AI development lifecycle, could include: ensuring acceptable outcome criteria are defined prior to model training; the management and disclosures of model and experiment asset inventories; AI and machine learning ("ML") corporate policies (similar to SDLC policies); and, ensemble modeling, which can help reduce risk via polling or combination-of-strengths mechanisms. Providing a framework that would be adaptive to the various risk tolerance levels and allow for

Olive
99 E Main Street
Columbus, OH 43215

higher levels of risk to be considered acceptable at appropriate stages of development, would promote transparency and innovation in AI development and result in better performing AI models.

5. **Managing Oversight & Accountability**

An AI Risk Management Framework should also promote and encourage collaboration and peer review within the developer community. Institutional Review Boards ("IRB") at academic institutions and hospitals impose a strict process for clinicians or scientists looking to conduct research. These researchers must submit an extensive application with a full proposal and study protocol and show evidence of completed training in human subjects research. Yet, similar oversight and governance does not exist for ML models in the U.S. There may be IRB requirements to begin the modeling but not for deployment. The researchers or vendors (if involved) may also have competing interests in deploying this work. There is no guidance on an industry standard for how parties should go about resolving these conflicts prior to deployment, if at all. Similarly, there are no checks and balances on the use of "off the shelf" AI tools and whether those are implemented in an appropriate manner. Any such peer review process should be proportionate to the purpose, proposed use, and risks presented.

As AI and ML are implemented and deployed, we are all learning. Mistakes are inevitable and lessons are being learned. Everyone does better if we can learn from one another's mistakes and not have to go through the pain of making the same one.

Any framework around AI governance, such as the AI Risk Management Framework, should prioritize and encourage transparency among developers, even on issues – such as bias – that are typically shielded from public view. This type of collaboration would also address the resource and knowledge gaps around AI. The ability to bring together resources would not only remove barriers to entry for market participants, but it would also allow market participants to expand access to computing resources that can provide better experimentation, training, validation, benchmarking, and evaluations of AI model and system metrics. In addition to the bias discussed above within AI models themselves, the knowledge gap creates another layer of bias given the limited number of experts and individuals researching and developing AI principles, model architectures, and frameworks.

There are examples of this type of collaboration in the cybersecurity community and vulnerability information sharing. Additionally, code coverage scores, which once were only seen by developers on a given project, are now universally expected on any public-facing or

Olive
99 E Main Street
Columbus, OH 43215

largely shared code repository. Collaboration within AI could be done through a similar centralized bulletin system that is promoted through the AI Risk Management Framework.

Conclusion

Olive appreciates the opportunity to respond to NIST's Request for Information on the AI Risk Management Framework and welcomes additional discussion on the development of industry guidelines within AI that will allow the AI community to continue to build meaningful, impactful, and smarter solutions. Please do not hesitate to reach out to Kelli Briggs, Head of Government & Industry Affairs, at kelli.briggs@oliveai.com.