# Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

September 20, 2021

Attn: Wendy Wigen
National Coordination Office (NCO)
2415 Eisenhower Avenue
Alexandria, Virginia 22314, US

RE: *Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource*

Dear Ms. Wigen,

Cerner Corporation appreciates the opportunity to submit public comment on the White House Office of Science and Technology Policy and National Science Foundation's RFI on an Implementation Plan for a National Artificial Intelligence Research Resource. As a leading supplier of clinical and management information systems we believe our experience provides us with valuable insight in this subject area and are grateful for the ability to share that insight.

If you have any questions or if we can provide any additional information, please do not hesitate to contact me at ███████████████

Sincerely,

*John Travis*
*Vice President & Regulatory Strategy Executive*
Cerner Corporation

First, Cerner would like to express our appreciation to the Office of Science and Technology Policy and National Science Foundation, and the National Artificial Intelligence Research Resource (NAIRR) Task Force, for their focus on this subject. We believe it is a critical topic for all industries – particularly as Artificial Intelligence (AI) and Machine Learning (ML) continue to become more established parts of research, data analytics, and software development. We believe the NAIRR Task Force's focus and investment with this Request for Information (RFI) and subsequent review is a critical step towards achieving a higher level of focus on AI and ML within the industries.

This RFI relating to AI research spans all industries. Cerner's responses to the RFI questions are from a healthcare industry perspective.

As various types of data contribute to a national AI database, it is critical to consider the sensitive nature of healthcare data, whether it be clinical or financial, compared to many other types of data. People may fear discrimination from an employer or from an insurance company based on disclosure or secondary use of certain health information. Needless to mention, the patient safety considerations and potential discriminatory uses of health information is heightened just by the mere fact that it is data related to human health.

It is for these reasons that the data (and NAIRR as a whole) should be segmented based on purpose and scope within the infrastructure. For example, for the reasons outlined above, healthcare data should be treated differently than other types of data. Access to healthcare data must adhere to Health Insurance Portability and Accountability Act (HIPAA) Privacy and Security rules, including de-identification via removal of Personal Health Information (PHI). If raw PHI is needed to compute a derived value, when the computation is completed, the PHI must be discarded.

1. **WHAT OPTIONS SHOULD THE TASK FORCE CONSIDER FOR ANY OF ROADMAP ELEMENTS A THROUGH I BELOW, AND WHY? [PLEASE TAKE CARE TO ANNOTATE YOUR RESPONSES TO THIS QUESTION BY INDICATING THE LETTER(S) OF THE ITEM (A THROUGH I IN THE LIST BELOW) FOR WHICH YOU ARE IDENTIFYING OPTIONS.]**

*B. A plan for ownership and administration of the National Artificial Intelligence Research Resource, including:*

> *i. An appropriate agency or organization responsible for the implementation, deployment and administration of the Research Resource; and*

**Cerner comments**
Unless AI ownership is segmented out in different categories by type and purpose of collection (e.g., healthcare, defense, etc.), it will likely be infeasible to have a single resource that can be fully shared and serve all unique needs.

*ii. A governance structure for the Research Resource, including oversight and decision-making authorities*

**Cerner comments**
It is critical that the governance structure for the NAIRR involve representation from knowledgeable and relevant public-private entities for the types of data and purposes of collection that prevail. This is necessary to ensure the requisite knowledge set is present to appropriately inform key decisions about the data, the development lifecycle of that data, AI/ML techniques, the implications of obtaining that data, and how the data should be handled related to privacy.

*D. Capabilities required to create and maintain a shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country, including provision of curated data sets, compute resources, educational tools and services, a user-interface portal, secure access control, resident expertise, and scalability of such infrastructure*

**Cerner comments**
Healthcare poses a unique use-case which necessitates special considerations such as elevated protections for PHI. The associated infrastructure strategy for the NAIRR must account for that.

2. **WHICH CAPABILITIES AND SERVICES (SEE, FOR EXAMPLE, ITEM D ABOVE) PROVIDED THROUGH THE NAIRR SHOULD BE PRIORITIZED?**

**Cerner comments**
Priority needs to be on publishing standards as part of this resource that provide the industry with guidelines for development and testing of AI software. As the FDA develops these guidelines for health information through ongoing guidance development for AI/ML, this is an opportunity to integrate with that organization to establish consistency between what they're providing and how this resource functions.

3. **HOW CAN THE NAIRR AND ITS COMPONENTS REINFORCE PRINCIPLES OF ETHICAL AND RESPONSIBLE RESEARCH AND DEVELOPMENT OF AI, SUCH AS THOSE CONCERNING ISSUES OF RACIAL AND GENDER EQUITY, FAIRNESS, BIAS, CIVIL RIGHTS, TRANSPARENCY, AND ACCOUNTABILITY?**

**Cerner comments**
While it is important to cultivate innovation, the principles of ethical and responsible research and development of AI cannot be neglected.  The code of conduct developed to provide guidance on how to abide by these principles should encourage self-regulation and self-reporting.  These principles fall into the following three categories:
- Data Responsibility
- Ethical Principles

- Transparency and Accountability

## Data Responsibility

Data policies and procedures should be documented, maintained, and enforced to:
- Ensure technical security and user security
- Provide system-level logging and security
- Provide incident management
- Provide change management
- Provide contingency planning for disaster recovery and resiliency
- Not use or disclose/share the data in ways other than stated, or as otherwise required by law
- Safeguard the data to prevent such misuse or unauthorized disclosures
  Report any misuse or unauthorized disclosure as soon as known
- Refrain from identifying or contacting subjects represented by or within the data

## Ethical Principles

Similar to principles developed by organizations like the Center for Practical Bioethics, the Berkman Klein Center at Harvard University the Stanford Institute for Human-Centered Artificial Intelligence, ethical principles must be applied to the use of the data.
- Literature research should be performed on the subject of interest with special attention on the topics of bias and fairness.
- Communication with Subject Matter Experts and other stakeholders is critical.
- Data sources must be understood and collection methods to identify potential opportunities for data bias must be examined.
- Diverse cohorts should be used in order to mitigate bias.
- Possible operational and usage patterns should be defined to identify any potential consequences of the AI's use that may affect some sensitive groups (e.g., racial, and ethnic minorities).
- Possible AI design approaches that may minimize bias should be identified.
- Fairness and bias should be measured using quantifiable, mathematical metrics.
- Expectations for transparency to users from the AI should be defined.
- Transparency of the AI development and validation process to consumers of the AI as it relates to potential for bias should be ensured.
- Analysis of appropriate performance and bias metrics for relevant subgroups should be performed. Typical subgroups include gender, race, ethnicity, age, marital status, location and time period.
- When a subgroup exhibits unacceptable performance that cannot be improved, the AI should be documented as not to be used on that particular subgroup.

As AI is used to drive toward positive outcomes, the most important aspect is to ensure the algorithms are monitored over time to retain quality and accuracy.

Quality management systems should be in place to ensure processes are followed, information is used in the same way, and methods are consistent.  This is not to say a national quality regulatory framework that requires submission thru a regulatory body or something along the lines of a 510(k) is needed, rather that quality management systems should be utilized as they exist for software today, similar to the <u>FDA's Quality System Regulation (QSR)</u> and the <u>International Organization for Standardization's (ISO) quality management system</u>.

If the monitoring identifies a need for change, end users must be notified, use of the intelligence must be discontinued if warranted until a change can be implemented and note of corrective actions must be taken. If the reason a change is needed is to prevent an adverse event or significant financial, regulatory or privacy burden, a process should be in place to determine root cause and create mitigation(s) to prevent such an issue from reoccurring.

Adverse events should be detected during monitoring by:
- Applying the AI to two different demographics.  Comparing the outcomes against both demographics can proactively assess how much risk is due to the demographics.
- Examining the AI results for an unexplained minority group effect that is statistically significant compared to the median.
- Explicitly creating a quality measure of disparity that can be monitored over time.  Changes in this quality measure can be investigated as new data is introduced.

Bias should be combated through responsible development practices and monitoring of AI in real world use. A monitoring framework and ecosystem should exist with the following capabilities to:

- Detect a worsening outcome disparity or performance bias. Mitigation options should include:
  - stopping use of the AI where it is performing poorly,
  - modifying the AI to correct the disparity or bias, and
  - proactively seeking out and considering users' feedback on their perception of bias.
- Regularly produce automated reports on the distribution of data inputs, performance and bias so they can be reviewed for:
  - changes in the mean or variance of a continuous data element,
  - changes in patterns of missing data,
  - changes in the distribution of categorical data – especially if a new category appears, as well as,

- o changes in results based on demographics factors.
- Produce near real-time alerting when a critical failure occurs.
- Proactively look for measures that may indicate a disparity in outcomes arising from the use of AI across demographic lines such as race, age, gender, or some combination of them.

### 4. WHAT BUILDING BLOCKS ALREADY EXIST FOR THE NAIRR, IN TERMS OF GOVERNMENT, ACADEMIC, OR PRIVATE-SECTOR ACTIVITIES, RESOURCES, AND SERVICES?

**Cerner comments**
No comments

### 5. WHAT ROLE SHOULD PUBLIC-PRIVATE PARTNERSHIPS PLAY IN THE NAIRR? WHAT EXEMPLARS COULD BE USED AS A MODEL?

**Cerner comments**
As mentioned above in question #1 under letter item B.ii, a governance board with representation from public-private participants is critical to contributing to oversight and decision-making about the data, understanding the data lifecycle, and protecting the privacy of the data (e.g., PHI, etc.) so that implications of decisions are well understood.

Accordingly, a distinction or separation should also exist that would delineate the management of the data assets from public policies/laws and the actual AI development. Public policies and laws are government core competencies. However, administration of the data asset and associated technical aspects may be more suited for a public-private partner.

### 6. WHERE DO YOU SEE LIMITATIONS IN THE ABILITY OF THE NAIRR TO DEMOCRATIZE ACCESS TO AI R&D? AND HOW COULD THESE LIMITATIONS BE OVERCOME?

**Cerner comments**
No comments