EMERGING TECHNOLOGIES ACCELERATOR (ETA) PROGRAM

AMENDMENT 0003

Q&A

Note: Q&A Window closed 13 May 2026.

74. MOVES: Technical Deliverable 1 specifies frequency response > 12Hz. Can you clarify what this means? Should systems be capable of detecting/reproducing tremors at frequencies up to 12 Hz, or must they handle frequencies greater than 12 Hz?

    Answer: Algorithms must be capable of reproducing tremors of at least 12 Hz.

75. MOVES: Is there a minimum or expected frame rate for videos?

    Answer: Videos will be standard framerates of 24 to 30 frames per second.

76. MOVES: For subjects with conditions in the training data, will there be 100 subjects per condition, or 100 subjects total?

    Answer: Target is roughly 100 subjects total.

77. MOVES: Will Algorithm 2 have access to the original RGB video as well as the outputs of Algorithm 1?

    Answer: Algorithm 2 can use the original RGB video.

78. MOVES: Is there a mechanism to consider rejection evidence? For example, the lowest ranked clips may contain evidence that rejects the possibility of a condition.

    Answer: Performers may use any information in the video to generate their diagnosis.

79. MOVES: Are conditions graded independently or together?

    Answer: The accuracy metric is overall top-1 accuracy.

80. MOVES: Is evaluation performed from a single view (unlike training data, which is collected from multiple vantage points)?

    Answer: Test data will be from a single view.

81. COSMIC: What type(s) of two-dimensional information to be predicted using the PPD are within scope?

    Answer: The Government is asking for two-dimensional imagery; however, this may change throughout the course of the program.

82. COSMIC: What types of existing GEOINT models does the PPD need to ingest?

    Answer: PPD inputs will vary throughout the course of the program. No PPD models are specified currently.

83. COSMIC: Will solutions be expected to ingest or emulate hyperspectral data? What about 3D GEOINT data such as LiDAR point clouds?

    Answer: Ingestion of the emulation of hyperspectral data may be required. Geolocated 3D data may also be an input into the PPD.

84. COSMIC: What can you tell us about the established change detection algorithms our solutions will be tested against? What is the level of precision for these algorithms?

    Answer: Detection algorithms will be provided at kick-off.

85. COSMIC: The Program Waypoints, Milestones, Deliverables section states there will be 9 evaluation milestones, but Table 2 only lists 6. How many evaluation milestones should we expect and when?

    Answer: Via Amendment 0003 (Tracked Changes), this has been corrected.

86. COSMIC: Is independent contractor data collection within scope?

    Answer: Yes, if it can be delivered to the Government.

87. COSMIC: If we use our own training data (collected independently, or obtained through other Government contracts), must we share it with all other Performers?

    Answer: Training data is expected to be delivered data, which may be shared with other Performers.

88. COSMIC: Will Performers have access to relevant agentic workflows to assist with the development and testing of our approaches?

    Answer: Example Agentic workflows will be provided at kick-off.

89. COSMIC: Can you please explain in more detail how the proposed metrics will be computed? For Change Detection, the target performance is 25% reduction in the time to detect change. Is time referring to the run-time of the processing system, or the amount of time that passes after a change event until it is detectable?

    Answer: The amount of time that passes after a change event until it is detectable.

90. Is an OT-CPFF (expenditure) contract type possible for an award on an ETA effort? The solicitation states that the Government anticipates cost reimbursement proposals; however, it also states that resultant OTs will reflect Fixed-Cost Milestones/Deliverables, suggesting a possible OT-FFP contract type. Are both contract types possible?

    Answer: Yes, it is possible. Reference ETA Program Announcement, Amendment 0002, Q&A # 73.

91. For possible OT-FFP contract type, are milestone payments intended to be tied to the listed Program Technical Deliverables, which include quantitative performance requirements?

    Answer: Payments will be tied to milestones/deliverables. Also reference ETA Program Announcement 6.B.2 Technical Data – Withholding of Payment.

92. The Technical Proposal outline includes some redundancy on some items – cover page included in Section 1.A. as well as Attachment 1 in appendix, and RDMP included in Section 2.L. as well as Attachment 6 in appendix. Can you please clarify?

    Answer: Via Amendment 0003 (Tracked Changes), this has been clarified.

93. LOCUS: The initial pre-solicitation refers to the absence of any metadata during the test & evaluation. Will there be additional information/data available during the inference for cross-referencing? If not, will Performers have access to the web (e.g., API calls) at runtime?

    Answer: Reference data may exist during test and evaluation. This will be emulative of reference data that will be available in operation. However, in operation web access will be restricted.

94. LOCUS: In indoor environments, geolocation ranges from difficult to technically impossible. Could you clarify if we are expected to address the entire spectrum, including "featureless" cases where cues are essentially absent? Alternatively, should we prioritize specific facets such as background noise and echoing, or environmental and cultural markers? Defining this scope could help us focus on the most relevant challenges.

Answer: Yes, Proposers are expected to address the entire spectrum.

95. LOCUS: Some modalities like video may have more anchoring information than a single image or audio clip. How will CE90 < 250m be used across different data inputs and how will aggregate metrics across instances from test-dataset be calculated e.g. percent of test-data that satisfies CE90<250. Also, will these metrics be further sub-grouped or bucketed e.g. indoor vs outdoor, image vs video

    Answer: Further information about test and evaluation strategy will be shared at kick-off.

96. LOCUS: Counter-intelligence/Adversarial attacks – is intentional noise/adversarial injection (e.g., corrupted/fake audio background, redacted visual cues etc.) in scope for T&E evaluation? Is the program interested in the aspect of improving the framework against these attacks?

    Answer: No, however ablated or redacted features are anticipated to be included in test events.

97. LOCUS: Can you provide the expected T&E runtime environment for container execution (CPU model/count, GPU type/count and memory, system RAM, local disk quota, OS/base image constraints), and whether GPU acceleration will be available for performers at test events?

    Answer: Details about the computational environment and computational restrictions will be provided at program kick-off.

98. LOCUS: Is it permissible to include locally hosted open-source/open-weight foundation models (e.g., LLMs or VLMs) packaged entirely within the Docker image as dependencies, provided the system is turnkey and runs fully autonomously? If yes, are there limits on model size, startup time, disk footprint, or licensing terms?

    Answer: Details about the computational environment and computational restrictions will be provided at program kick-off. Systems may be expected to run on standalone networks.

    According to the ETA Program Announcement:

    IARPA requires sufficient rights to IP developed or used in the conduct of the proposed research to ensure that IARPA can successfully (a) manage the program and evaluate the technical output and deliverables, (b) communicate program information across Government organizations, and <u>(c) support transition to and further use and development of the program results by IC users and others.</u>

99. We have submitted our white paper. Is feedback anticipated on a rolling basis, or at a specific time after the BAA has closed?

    Answer: After 27 May 2026, on a rolling basis.

100. MOVES: What type, approximate size, and timing should proposers expect for the training dataset that the T&E team will collect and distribute to performers?

Answer: We expect roughly one hour of total video from each unique volunteer performing activities of daily living. That will be spread across multiple separate video files whose lengths may be from 30 seconds to 20 minutes. These targets may be adjusted if experimental conditions require a change.

101. MOVES requires performers to select three conditions from the listed set: Parkinson's Disease, Essential Tremor, Dystonia, Stroke, Mild Traumatic Brain Injury, and Multiple Sclerosis. Should performers identify the three selected conditions in the white paper, and may the Government direct a change in selected conditions after award, based on dataset availability?

Answer: Performers should specify the selected conditions in their White Papers. Performers may change, or IARPA may require a change, if availability of data makes that the best path forward.

102. MOVES: What documentation will be required for the use of public/open-source videos, clinical videos, or other human-subject video data in development or training? Should proposers assume that any collection of new human video data will require IRB approval and IARPA HSR review before use?

Answer: The use of public video or open-source datasets must be reported to IARPA. Clinical or other non-public video data containing people the Proposers may have or collect is human subjects data and requires IRB approval as described in the ETA Program Announcement.

103. MOVES: Technical Deliverable 1, the program describes the output of skeletal movement and simulated unclothed full-body animation. Is the intended output a diagnostic kinematic/body-shape abstraction for clinician review rather than a photorealistic anatomical reconstruction? Are there specific privacy, identifiability, or fidelity constraints for this output?

Answer: The unclothed full-body animation is intended to illustrate the motion of the body unobstructed by clothing, for clinician review in the context of diagnosing movement disorders. It does not need to be photorealistic, and identifiability of the real individual from the unclothed body simulation is not a concern.

104. Will the Government provide a required interface or file format for MOVES keypoint outputs, skeletal animations, body-shape outputs, diagnostic segment rankings, and diagnosis scores?

Answer: IARPA and the T&E team, taking into account the technical approaches of the selected Performers, will specify the interface and file format requirements after the start of the program.

105. MOVES states that outputs from one team's Technical Deliverable 1 may be used as inputs for another team's Technical Deliverable 2 during T&E. Will performers be required to design Technical Deliverable 2 to accept a common standardized TD1 output format independent of their own TD1 implementation?

     Answer: IARPA and the T&E team, taking into account the technical approaches of the selected performers and evaluating the costs and benefits of doing so, may require a Team's Technical Deliverable 2 to conform to common interface standard.

106. For the MOVES diagnostic-accuracy metric, will the 80% target be assessed at the subject level, video level, or segment level? Will cases be balanced by subject, video, or 15-second segment?

     Answer: Metrics will be assessed at the unique subject level. Best efforts will be made to balance cases by subject.

107. Will MOVES ground truth labels be based on confirmed clinical diagnosis, clinician rating of video-observable symptoms, or another adjudication process?

     Answer: Ground truth for the training dataset collected by the T&E team will be based on clinical diagnoses. The test set may include open-source video with clinician rating of video-observable symptoms.

108. LOCUS: We would like to know whether the reference datasets used by various teams will be available to ALL teams participating in the LocUS program. For example, if a team used part of their budget to purchase a particular reference dataset, would it become available to all other teams?

     Answer: We cannot guarantee the release of all reference datasets collected by other teams.

109. ARCADE: We are writing to request clarification on the scope of the source code delivery requirement in the ARCADE Software Deliverable Formatting section of Attachment 1 (page 6) to the ETA Program Announcement.

     A. Does "source code" in the Software Deliverable Formatting section refer to the complete application source code, or to the build scripts, configuration, and deployment code necessary to reconstruct and modify the containerized deliverable for T&E deployment constraints?

        Answer: The government prefers to have complete application source code, but at a minimum the build scripts and configurations necessary for T&E to run test and evaluation are required.

     B. Will the Government accept a proposal that includes a separate option price for delivery of source code associated with commercial or privately funded software? The source code for such software will not be produced via the OT nor funded exclusively or in part by the Government. What is the Government's rationale for requiring delivery of privately funded IP?

Answer: According to the ETA Program Announcement, which applies to all ETAs:

IARPA requires sufficient rights to IP developed or used in the conduct of the proposed research to ensure that IARPA can successfully (a) manage the program and evaluate the technical output and deliverables, (b) communicate program information across Government organizations, and (c) support transition to and further use and development of the program results by IC users and others. IARPA can achieve this mission through Unlimited Rights. There may be any number of approaches to IP rights to achieve IARPA's program goals. In addressing the approach to IP rights, Proposers shall (1) describe the intended use of patented invention(s) or data, including, technical data and computer software, in the conduct of the proposed research; (2) describe the rights being offered to the Government along with a justification if less than Unlimited Rights is being offered; (3) explain how IARPA will be able to reach its program goals with the rights offered to the Government; (4) identify the cost to the Government to acquire additional or alternative rights beyond those being offered, if applicable; and (5) provide possible alternatives in any area in which the offered rights may be insufficient for IARPA to achieve its program goals (e.g., the possibility of future licensing of privately-developed software to U.S. Government agencies at a reasonable cost.)

C. Will the Government amend the solicitation to allow a performer's proposal to include an alternative approach to enabling modification of Performer containers to suit deployment constraints? The concern is that requiring delivery of privately funded source code is restrictive and could limit competition, as a proposal that provides an effective means for the Government to modify Performer containers but does not include delivery of source code could be determined non-responsive.

Answer: No. Reference answer to Question # 109B, immediately above.

110. DECIPHER: In the section on metrics, under "Detection", this sentence appears:

"Performance in retrieving documents from a target domain of interest may also be measured relative to a baseline, which may be a general-purpose, state-of-the-art information retrieval capability."

The detection TA1 task is described as identifying words and phrases in a text that are specialized language, not in performing an information retrieval-style document identification task. Is there an additional information retrieval task that is wanted in TA1 or TA2, and if so, what is the task?

Answer: Further information about TA-1and TA-2 tasks and metrics will be shared during the program.

111. DECIPHER: Language Variety: Should the proposed solution be effective across many different languages (e.g., 10+ high-resource languages), or would demonstrating effectiveness in a smaller set of languages be sufficient for this 12-month program?

Answer: Specific medium and high resource languages of interest in the program will be shared at kick-off. Demonstrating effectiveness in a smaller set would be sufficient; however, we are interested in solutions that apply to many different languages.

112. DECIPHER: Compute Hardware: May we budget for one GPU workstation owned by the prime (accessible to the full team) and one GPU-equipped laptop for deployment testing, consistent with program requirements?

Answer: Specialized hardware to achieve program goals is permissible. All equipment purchased for the program is considered Government owned.

113. Regulatory Framework: Since this is an OTA effort, could you confirm which regulatory framework applies to Performers (e.g., FAR 52.245-1, 2 CFR 200, or OTA-specific terms)?

Answer: For this opportunity, the FAR does not apply. This is also not Financial Assistance (e.g. Grant or Cooperative Agreement). The selected Proposer will eventually receive a copy of the Other Transactions Agreement (OT) containing the Terms and Conditions (Articles) to negotiate with the Agreements Officer.

114. DECIPHER: Data Licensing: Q&A # 64 indicates that proprietary or licensed data is acceptable if the Government and T&E team have full internal use rights, even if public redistribution is restricted. Could you confirm that such a license remains in scope for DECIPHER?

Answer: Yes

115. What is the anticipated number of awards?

Answer: Unknown at this time.

116. What is the anticipated funding per award?

Answer: Unknown at this time.

117. DECIPHER: Are proposals that use program resources to develop annotated datasets for development purposes in-scope?

Answer: Yes.

118. DECIPHER: Are technical approaches that require some amount of language-specific configuration (e.g. uploading of a background text corpus, designation of language-specific NLP processors) in-scope? If so, will T&E provide information to Performers about the natural languages represented in the datasets used for evaluation?

     Answer: Details will be provided in advance of the evaluations, where they are necessary.

119. DECIPHER: Will the test sets used by the T&E team to evaluate performer systems consist of authentic or synthetically generated datapoints?

     Answer: Details on the test events will be shared at kick-off.

120. DECIPHER: Does the government have an opinion on whether synthetic data is sufficient for training, validating, and/or evaluating DECIPHER performer systems?

     Answer: All data approaches will be considered.

121. DECIPHER: Is the government open to proposals that, due to IP considerations surrounding proprietary data used for training models, do not deliver training data to the Government, but do give the Government, Government Purpose Rights to the final model weights?

     Answer: Reference answer to Question # 109B above.

122. DECIPHER: Is it desirable for systems to make use of metadata in addition to the natural language content of datapoints? If so, will the Government provide a standard format for how metadata will be structured in evaluation datasets used by T&E, as well as a list of what metadata fields can be expected to be available?

     Answer: Additional details will be provided throughout the program.

123. DECIPHER: Can we pursue research questions that will open up promising new lines of inquiry, but will only partly inform the airgapped capability at the end of 1 year of POP?

     Answer: Strong White Papers/proposals directly address the requirements in the announcement.

124. DECIPHER: Is there any interest in making some portion of the program outputs open source and if so, will IARPA consider a similar model to the DARPA Open Catalog from XDATA and MEMEX?

     Answer: We will take this into consideration.

125. DECIPHER: What natural languages will be represented in the T&E evaluation datasets? If not known at this time, will T&E inform performers of the natural languages present in the datasets in advance of evaluation events?

Answer: Further information about evaluations will be provided during the program.

126. DECIPHER: Does IARPA have a greater interest in slang/technical jargon or obfuscated/coded language?

Answer: No.

127. DECIPHER: Are proposals that take different approaches to slang/technical jargon vs. obfuscated/coded language in scope?

Answer: Yes

128. Can the Government provide a format, template, content requirements, and/or slide count for the (optional) in-person presentation? Without such guidance, it is challenging to start preparing presentation material ahead of time. Although this is an option that may or may not be exercised by the Government after white paper evaluations, any information that can be provided ahead of time will improve the quality of vendor responses.

Answer: Additional information will be provided later with plenty of time to prepare.

129. Can the Government provide information on how much time vendors will be allowed between the invitation to present in person and the meeting date itself (or the due date for oral slide content if required in advance)?

Answer: Additional information will be provided later with plenty of time to prepare.

130. Can the Government provide information on how much time vendors will be allowed to prepare a written proposal after notification of white paper viability?

Answer: Additional information will be provided later with plenty of time to prepare.

131. MOVES: The algorithm requirements on Page 32 state that the algorithms "allow the user to select a subject of interest." Can IARPA provide additional details on how that will occur? Is the expectation that Performers will develop and deliver a GUI for a user to select a subject from a displayed video? Or will the T&E team create a standardized UI with an Interface Control Document (ICD) for interoperability with Performer algorithms?

Answer: The intent is the Performer will provide the simplest method of selecting a subject of interest, which may include a command line or GUI. This will be coordinated with the T&E team after kickoff.

132. MOVES: Please confirm whether the intended dataset for training and testing purposes will include 100 subjects for each neurological condition or 100 subjects total across all conditions?

Answer: The target is100 total subjects.

133. MOVES: Can further details be provided on the data variety of the planned Government provided dataset? Will video data be provided for each subject that includes multiple viewpoints, multiple clothing outfits, multiple environmental/lighting conditions? Can we expect that some subjects will be imaged with different clothing but the same medical condition? Subjects with similar clothing but different medical conditions? etc...

Answer: Best efforts will be made to provide all this information.

134. MOVES: For Algorithm 2, please clarify if a user would indicate the specific neurological condition to diagnose at algorithm execution time OR if the algorithm is to diagnose an input video for each of the 3+ performer-selected conditions simultaneously with separate likelihood scores for each of the conditions reported?

Answer: The user will not specify a condition.

135. COSMIC: Are performers allowed to build on top of an existing geospatial AI model (such as Clay) by fine-tuning or adapting it for this effort, or does that fall under the "no new foundation models" restriction in the solicitation?

Answer: Yes, it is acceptable.

136. COSMIC: Will the Government supply the change detection algorithms that will be used to test our solution, or are performers responsible for providing those?

Answer: Example change detection algorithms will be provided at kick-off.

137. COSMIC: What data will the Government provide to Performers, and roughly how much? Also, if our team needs to purchase commercial satellite imagery or other data for training purposes, is that an allowable cost under the award?

Answer: The volume of data provided is best described as exemplary data vs. quantities required for training. Commercial satellite imagery or other data for training purposes is an allowable cost as long as licensing it allows it to be delivered to the Government.

138. COSMIC: The solicitation requires the solution to run on standalone systems. Can the Government clarify what that means in practice? For example, are we talking about disconnected networks, classified environments, limited hardware, or something else?

Answer: Disconnected networks with hardware limitations that may be classified.

139. COSMIC: The solicitation states that development of “novel Agentic AI systems” is out of scope. Can the Government clarify whether Performers are allowed to develop custom workflows, orchestration layers, tools, or interaction methods around existing LLMs/agentic systems for querying and interacting with the PPD?

    Answer: Yes, this is allowable.

140. LOCUS: Will performers be expected to share data?

    Answer: Performers are expected to deliver their data to the Government.

141. LOCUS: Is system output expected to be limited to point-prediction? (e.g., Some imagery may only be locatable to a region or country)

    Answer: Further information about system output expectations will be shared at kick-off.

142. LOCUS: Will altered or edited inputs be considered for queries? (e.g., AI generated/inpainting, blurring license plates / faces)

    Answer: Yes.

143. LOCUS: How will explainability of the system be quantified?

    Answer: The explainability assessment will be discussed during kick-off.

144. LOCUS: Will test queries include multi-channel audio (stereo, binaural, or higher-order ambisonic), and if so, can Performers expect a known channel layout? If queries may be either mono or multi-channel, will systems be evaluated on the same metric across both?

    Answer: Mono and stereo audio will be included and will be evaluated on the same metric.

145. LOCUS: What fraction of test queries are expected to be indoor vs. outdoor?

    Answer: Composition of test queries may not be shared with Performers. However, after test events, we will provide information about performance in different conditions.

146. LOCUS: Is the 5-minute limit per-query or test-set average?

    Answer: Details of test and evaluation processes will be provided at program kick-off.

147. LOCUS: Will integration / inference cloud platform be provided as part of the program for testing and evaluation?

    Answer: No. Performers are expected to perform their own integration into the test and evaluation cloud platform. ODCs for specialized hardware for core development are available.

148. LOCUS: Will performers be provided access to any tools? (e.g., Gemini, GPT) Will credits/tokens be provided by the program?

    Answer: No. However, costs for tools should be included in your proposal.

149. LOCUS: Is the system expected to operate at full performance in air-gapped environments (with no internet access)?

Answer: Yes.

150. ARCADE: Is the expectation that the solution will provide recommendations for a single component or groups of components?

Answer: Multiple components are acceptable with reasoning and comparison analysis behind them.

151. ARCADE: Is the expectation that the solution will recommend schematic connectivity?

Answer: No. Schematic creation is out of scope for ARCADE. However, the Government is open to hearing about creative solutions in this area.

152. ARCADE: When will final program metrics be provided?

Answer: Program metrics will be provided at kick-off and adjusted throughout the program.

153. ARCADE: Can a subcontractor serve as the Principal Investigator?

Answer: It depends. Reference ETA Program Announcement, Section 2, 50 U.S.C. § 3024(m)(6).

154. The solicitation requests that offerors indicate their preferred type of OT agreement, mentioning CPFF and FFP as options. Given that OT agreements are not governed by FAR guidelines, could the Government please clarify the intended structure and characteristics of a "CPFF" agreement in this OT context?

Answer: This answer assumes your White Paper was determined viable and your company was invited to submit Volume I and subsequently Volume II.

Although the Proposer may choose a FFP Proposal, ETA Program Announcement, Section 4.B.2. states:

> The AO may consider granting the request for an FFP award under the following circumstances:
>
> - Small Businesses,
> - Proposers without an approved accounting system,
> - Start-Up companies,
> - Consortia that may include universities and non-profits or;
> - Where cost-sharing or government participation in the work is appropriate.

The Proposer, after reading ETA Program Announcement Attachment 1, will recommend Milestones/Deliverables in Volume I. During the Volume I evaluation, the Government PM may recommend changes or ask questions that may prompt the Proposer to make adjustments.

Once agreed upon, in Volume II, the Proposer will assign monetary values which will be fully supported by the proposed cost. Once negotiated and agreed upon, the OT will reflect the Fixed-Cost Milestones/Deliverables.

155. For offerors proposing an FFP structure, would the Government consider accepting a high-level cost summary (in the cost proposal) in lieu of the detailed cost element breakdown typically associated with cost-reimbursement proposals?

Answer: No, reference answer to Question # 156 immediately above. Also, ETA Program Announcement, Section 4.B.2. states:

> Proposers requesting other than cost reimbursement (i.e., FFP) may be directed by the AO to provide "other than certified cost or pricing data" and/or cost supporting information.
>
> The cost proposal shall contain sufficient factual information to establish the Proposer's understanding of the project, the perception of project risks, the ability to organize and perform the work and to support the realism and reasonableness of the proposed cost, to the extent appropriate.

156. For OT agreements structured similarly to a CPFF arrangement, are discrete priced milestones required, or will the agreement function more like a traditional FAR-based CPFF contract with deliverables and progress payments?

Answer: Reference Answer to Questions # 156 and 157 above. Not sure what is meant by "the agreement function more like a traditional FAR-based CPFF contract." The FAR does not apply.

157. If milestones are required under an OT-CPFF contract, will the contractor be authorized to invoice for the full milestone amount upon completion and acceptance of each milestone? If not, is the contractor expected to invoice only for actual costs incurred in achieving each milestone?

Answer: Reference Answer to Questions # 156, 157, and 158 above. Once a fair and reasonable Fixed-Cost Milestones/Deliverables table is agreed upon and finalized in the bilaterally signed OT, the <u>Performer</u> will invoice for the full amount.

158. If the OTs reflect a hybrid type agreement, could the Government provide further clarification on how this structure will be implemented within a cost-reimbursement framework, particularly regarding the interplay between cost reimbursement and fixed-price elements?

Answer: Reference Answer to Questions # 156, 157, and 158 above. Not sure what is meant by "hybrid type agreement." The resultant OTs will reflect Fixed-Cost Milestones/Deliverables.

159. MOVES: Will full camera matrices (or intrinsic and extrinsic camera parameters) be provided for training data?

Answer: No.

160. MOVES: Will IMU data included in training data contain precise hand motion information?

Answer: The IMUs will be on major joints.

161. MOVES: Which skeletal keypoints will be included in the algorithm 1 metrics? (Fingers? Feet? Ankles? Facial keypoints?)

Answer: Final details will be discussed at kick-off.

162. MOVES: Is there a limit to compute time?

Answer: No.

163. ARCADE: Will there be datasets or datasets provided to train the LLM? Or, the training needs to be performed using the datasheets or data available from the internet?

Answer: No. Datasets will not be provided to train LLMs.

164. ARCADE: Will there be a library of components provided?

Answer: Yes. An initial library of components will be provided at kick-off.

165. ARCADE: Will there be a benchmark provided to test the results?

Answer: Yes.

166. ARCADE: Can the Government clarify whether proposers must address the full end-to-end workflow across all identified task areas, or if focused proposals targeting a subset of task areas are acceptable?

Answer: Proposers must address all three task areas.

167. ARCADE: Can the Government provide additional guidance regarding the anticipated composition, scale, and classification level of source materials to be ingested (e.g., commercial technical documentation, structured/unstructured datasets, multimodal artifacts, etc.)?

Answer: ARCADE is an UNCLASSIFIED program. The library will include a wide variety of commercially available components, including passive, active,

electromechanical, RF, power-management, sensors, connectors, microcontroller families, and some FPGAs.

168. ARCADE: Does the Government anticipate any restrictions or preferences regarding the use of commercially available foundation models, externally hosted inference APIs, open-weight models, or hybrid architectures across anticipated deployment environments and classification levels?

    Answer: No preferences. Reference answer to Question # 109B above.

169. ARCADE: Can the Government provide additional guidance regarding anticipated deployment environments, including whether solutions are expected to operate within disconnected, air-gapped, cross-domain, or multi-classification environments?

    Answer: Deployment environments vary across government transition partners. ARCADE performers are encouraged to utilize LLMs during development. Final architectures should remain adaptable for specific customers requiring self-hosted, closed network deployments, but T&E will only evaluate solutions on the models chosen by performers.

170. ARCADE: Will the Government provide additional guidance regarding the anticipated evaluation approach, representative mission workflows, or operational scenarios that Proposers should align to within the white paper response?

    Answer: ARCADE performers are encouraged to utilize LLMs during development and provide complete codebase and documentation for the ARCADE T&E to evaluate using the same LLMs. Final architecture should remain adaptable for customers requiring self-hosted, closed network deployments, but T&E will only evaluate solutions on the models chosen by performers.

171. ARCADE: Beyond technical performance, what operational transition considerations will carry the greatest weight in evaluating prototype viability and downstream adoption potential?

    Answer: Transition requirements will be discussed throughout the program.

172. For the 1/3 cost-share criterion, may in-kind contributions (existing IP licenses, facility access, dedicated equipment, donated labor) count toward the 1/3 non-Federal share, or must it be cash outlay? If in-kind is allowed, what valuation methodology is acceptable?

    Answer: Yes, in-kind contributions are allowed, within reason. If the Government paid for the item, equipment, facility (fully or partially), or license via a separate Contract/OT, then no. The AO will base value from the submitted, verifiable information submitted by the Proposer. In your question, donated labor is best.

173. What is the anticipated turnaround from (a) white paper submission to viability notification, and (b) Volume I submission to selectability decision?

    Answer: Additional information will be provided later with plenty of time to prepare.

174. Where a small business prime intends to build a prototype on top of open-source AI components (e.g., Llama community license, Apache 2.0 models, MIT-licensed libraries) and concurrently uses related tooling on commercial engagements outside the IC mission, can the Government clarify (a) whether those upstream community licenses are compatible with the Unlimited Rights or GPR posture sought on deliverables, and (b) whether the Open Source assertion under Appendix A.3 is the appropriate vehicle to disclose those dependencies and any flow-down obligations?

    Answer: Yes, use Appendix A3 IP Rights. Also, reference answer to Question # 109B above.

175. If a deliverable contains commercial pre-existing software licensed to the Performer with use restrictions, does the GPR floor for mixed-funding deliverables apply only to new development, with the underlying commercial component retaining its commercial license terms?

    Answer: Reference answer to Question # 109B above.

176. For Proposers requesting an FFP award (small businesses, start-ups, no approved accounting system), are the indirect rate disclosures still required as supporting information, or simplified?

    Answer: Reference answer to Questions # 156 and 157 above.

177. Cost sharing is "encouraged where there is a reasonable probability of a potential application." How will the Government weigh the credibility of proposed commercialization paths during evaluation? Would prior agency awards in adjacent areas (e.g., DOE prize awards, NSF SBIR Phase I/II, prior IARPA selections) constitute supporting evidence?

    Answer: It is up to the Proposer to explain. It is not an evaluation criterion.

178. Where a Proposer's team includes a consultant who concurrently holds a separate, unrelated contract with one of the SETA support entities listed in Section 4.A.2, is this a disclosable OCI requiring mitigation?

    Answer: Yes.

179. COSMIC: Has IARPA established commercial data-access arrangements that Performers can leverage, or must each Performer secure independent commercial imagery licenses? If the latter, may license costs be charged as ODCs?

    Answer: IARPA may be able to provide some commercial data access and/or commercial data. If a Performer elects to secure an independent license, it must allow delivery to the Government. Yes, it can be included as an ODC in the Cost Volume.

180. COSMIC: Performance is measured against "existing change detection algorithms." Will the specific algorithms be enumerated at kickoff, or is the PPD framework expected to be detector-agnostic?

    Answer: The PPD framework is expected to be detector-agnostic; however, example detection algorithms will be provided at kick-off.

181. COSMIC: Will the Government specify the agentic AI systems against which the PPD framework must demonstrate compatibility?

Answer: Example agentic AI systems will be provided at kick-off.

182. COSMIC: Geolocation Error target is CE90 ≤ 10m - For SAR/EO harmonization, is this measured pre- or post-orthorectification, and which DEM or reference frame is presumed?

Answer: Post orthorectification. No DEM is presumed. Evaluation will be relative to a reference hold out image or other temporal outputs of the PPD system. Intent is to produce spatially consistent PPD.

183. COSMIC: What types of data may the system have access to generally? In the example of getting a count of tanks in Ukraine, is it expected that the system has access to a detection layer of information from a database? What is the expectation of incorporating that detection layer into a datacube?

Answer: The datacube should be agnostic to the detection process.

184. COSMIC: It is stated in the document "Once the Government has received sufficient White Papers per Program, the AO will issue an ETA Program Announcement Amendment re-labeling the Program Title in Attachment 1 – ETA Programs from "Active" to "Not Active. Could this happen before 30 June 2026?

Answer: Yes. The window to submit White Papers is guaranteed to stay open until 27 May 2026. This is applicable to all ETAs.

185. COSMIC: Would any of the downstream tasking processes be available to the PPD generation pipeline? For example, if there was a SAR->EO generative model in a process toolbox outside of this pipeline, could the PPD generation pipeline leverage it?

Answer: While downstream tasking processes may be used for the PPD generation, these processes will not be provided by the Government.

186. COSMIC: Will cloud compute resources be provided on the program?

Answer: Yes, some cloud computing resources will be provided by the Government. Specific hardware to train, develop, and test hardware is an allowable ODC to be included in the Cost Volume.

187. COSMIC: Has the government finalized details around performer access to agentic AI capabilities?

Answer: It will be provided at kick-off.

188. COSMIC: Will the government team be providing a test harness and/or evaluation data to the Performers so that they can evaluate potential solutions at a faster cadence? If so, how soon after work starts will T&E deliver the test harness and/or evaluation data?

Answer: Some data may be provided; however, the volume and delivery dates are TBD.

189. Can the government confirm that a minimum "Active" period will be maintained until at least June 15, 2026 to ensure offerors have enough time to prepare high quality submissions aligned with the Government's mission needs?

Answer: No. The window to submit White Papers is guaranteed to stay open until 27 May 2026.

190. DECIPHER: ETA Program Announcement, Attachment 1. Does the second column refer to the month within the program of each milestone?

Answer: Unless otherwise stated, it refers to the month of the period of performance. This is applicable to all ETAs.

191. DECIPHER: ETA Program Announcement, Attachment 1. Can the Government confirm that the intent of the scope requirement to be able to "detect and generate definitions for novel future terms" is for current proposed methods to be applicable to future specialized language terms as they emerge - versus being able to predict future coded language/jargon before it is used?

Answer: Yes

192. DECIPHER: ETA Program Announcement, Attachment 1. What is the expectation for a difference between Performers' systems and a current state-of-the-art information retrieval system in retrieving documents from a target domain of interest?

Answer: The Government does not understand this question.

193. ARCADE: ETA Program Announcement, Attachment 1. Can the Government clarify any difference in content or expectations between Phase 1 and Phase 2, or are those distinctions entirely based on Performers' development schedules?

Answer: Based on performance in Phase 1, Phase 2 metrics may be re-evaluated.

194. ARCADE: ETA Program Announcement, Attachment 1. The software deliverable formatting requires that the entirety of the performer's system be included within the delivered container. Since performers are encouraged to use state-of-the-art LLMs, with a final architecture that enables self-hosting, and even self-hosted LLMs are typically shared resources, may the container be delivered to call out to an LLM, that may be self-hosted or SaaS? Alternatively, must the container include a self-hosted LLM?

Answer: If the solution requires a self-hosted LLM, it should be included in the deliverable. Also, reference answer to Question # 109B above.

195. ARCADE: ETA Program Announcement, Attachment 1. Will T&E consider performance with a self-hosted LLM, with the Performer's preferred LLM, or both?

Answer: T&E will only consider the performance of the performers preferred LLM.

196. DECIPHER: Can you kindly provide details or an example of the prose description of the domain of interest that may be provided at test time?

Answer: Additional information will be provided at kickoff.

197. DECIPHER: The program description mentions "Performance in retrieving documents from a target domain of interest may also be measured relative to a baseline, which may be a general-purpose, state-of-the-art information retrieval capability." Is the intention that DECIPHER TA1 includes creating an information retrieval capability rather than or in addition to specialized language expression/phrase span detection?

    Answer: No.

198. Follow-on Production/Transition Criteria: Can the Government provide additional guidance regarding prototype performance, transition, or operational criteria anticipated to support eligibility for potential follow-on production activities?

    Answer: Other than what is included in Section 2 “Award Information,” there is no additional guidance at this time.

199. Teaming Flexibility: May performers refine or expand teaming composition between the White Paper and any invited proposal stage, provided the core technical approach remains materially consistent?

    Answer: Yes, but within reason. If a team member/partner who offered a key aspect of the technical solution leading to selection is no longer present, it may render the White Paper no longer viable. This is especially important if the member/partner who is no longer present was the NDC participating to a significant extent.

200. COSMIC: Evaluation Architecture/Interchange Expectations. Does the Government anticipate defining a common interchange schema, evaluation API, or standardized data exchange mechanism across COSMIC performers for T&E interoperability purposes?

    Answer: Yes, details will be provided at kick-off.

201. COSMIC: Pseudo-Persistent Data (PPD) Characterization. Is pseudo-persistence anticipated to be evaluated primarily through temporal continuity, operational utility, reconstruction fidelity, or a combination thereof?

    Answer: The PPD will be evaluated through a combination of reconstruction fidelity and operational utility.

202. COSMIC: Operational Utility vs. Reconstruction Metrics. Will COSMIC evaluation emphasis include downstream analytic effectiveness and operational utility in addition to image-centric reconstruction metrics?

    Answer: Yes.

203. COSMIC: Temporal Cadence Expectations. Can the Government provide additional guidance regarding anticipated temporal update cadence expectations for pseudo-persistent data generation (e.g., near-real-time, daily, weekly), and whether cadence expectations may vary by modality or area of interest?

    Answer: Temporal Cadence Expectations will vary during the lifetime of the program.

204. COSMIC: Evaluation / Deployment Environment Constraints. Can the Government provide additional guidance regarding anticipated deployment and evaluation constraints for performer-delivered COSMIC software, including expected compute environments, GPU availability, storage assumptions, and whether evaluations may occur in disconnected or air-gapped environments?

     Answer: We anticipate evaluation on stand-alone networks. Specific constraints will be provided at program kick-off.

205. Background IP Protection. Will the Government confirm that Government Purpose Rights associated with prototype deliverables do not extend to clearly identified background IP, pre-existing architectural methodologies, or separately developed frameworks identified within rights assertions?

     Answer: Reference answer to Question # 109B above.

206. Commercial Data Licensing Survivability. Can the Government provide additional guidance regarding acceptable approaches for commercial data licensing where redistribution rights may differ between development, testing, evaluation, and transition environments?

     Answer: Reference answer to Question # 209 immediately above.

207. If an NDC Prime performs the majority of work, does the inclusion of a traditional defense contractor as a subcontractor materially affect NDC qualification status for the prototype effort?

     Answer: The NDC must participate to a significant extent.

208. LOCUS: Can I include some preliminary results involving videos as a link in the white paper?

     Answer: Yes.

209. LOCUS: Can we use a vehicle license plate information as one of the features?

     Answer: No, images will have PII information removed such as faces and license plates.

210. LOCUS: Can we use human speech as one of the features?

     Answer: Yes.

211. The solicitation states Volume I is limited to 15 pages, attachments do not count, the transmittal letter is not to exceed one page, and the TDD section requires a Gantt chart that may use multiple pages if needed. Will the Government please confirm whether the 15-page Volume I limit includes the transmittal letter, cover sheet, and Gantt chart, or only Section 2 Detailed Proposal Information? Please also confirm whether a multi-page Gantt chart is permitted within the 15-page limit if needed for legibility?

    Answer: Via Amendment 0003, the Cover Sheet is now Attachment 1 only, excluded from the 15-page limit. Amendment 0003 has also moved the contents of the RDMP to Attachment 6. Sections 1 and 2 are included in the 15-page limit.

212. Volume I asks for cost estimates by cost element, while Volume II contains the detailed cost proposal and supporting BOE requirements. Because Volume II is only required if requested, but Volume I requires cost, schedule, and milestones, including cost estimates by cost element, please clarify the expected level of detail for Volume I cost estimates. Are labor categories, rates, subcontractor BOEs, and indirect-rate support expected in Volume I, or only a ROM by cost element and phase?

    Answer: For Volume I, ROMs by cost element for the PoP/Phase, Option Period(s) as applicable and total will suffice.

213. The solicitation says additional supporting/background materials not specified in Section 4 will not be evaluated. May proposers include compliance matrices, resumes, or additional corporate capability materials as attachments, or will such materials be treated as non-evaluated supporting/background materials unless explicitly listed in Section 4?

    Answer: With the exception of resumes, which can be included in Volume I, Attachment 8, the rest are unnecessary.

214. The announcement states the AO will request research-security documentation after Volume I selection, but also states IARPA requires completed NSPM-33 disclosure forms with each Volume I proposal. Will the Government please clarify whether NSPM-33 biographical sketch and current/pending support disclosure forms for covered individuals, and covered-institution certifications where applicable, are required with the initial Volume I proposal submission or only after the proposal is determined selectable/selectable with modification?

    Answer: Reference Section titled “Covered Individuals” where it states:

    IARPA requires the completed NSPM-33 disclosure forms to be submitted *with each proposal (Volume I)*.

215. The announcement defines covered individuals as key personnel who substantively contribute and are employed or compensated by the proposer or subcontractors, but examples include a broad set of paid research staff. Will the Government please confirm whether ‘covered individuals’ include only Key Personnel identified in the proposal or also significant contributors, consultants, and subcontractor staff below the Key Personnel level?

Answer: Key personnel identified as such in the Proposal.

216. MOVES: Regarding the T&E-collected training dataset of approximately 100 volunteers with movement disorders and approximately 100 volunteers without known movement disorders, please clarify:

    a. Will performers receive the IMU data?

       Answer: Yes.

    b. Will IMU motion capture be performed on both groups, or only on the movement disorder group?

       Answer: Yes, it will be performed on both groups.

    c. How many IMU sensors will be used per subject, what anatomical placements will be used, and will the placement protocol be made available to performer teams?

       Answer: Final details will be determined when the data is collected.

    d. What frame rate will the videos be (e.g., 30 Hz or 60 Hz)?

       Answer: Final details will be determined when the data is collected.

217. MOVES: The solicitation states T&E may feed one team's Deliverable 1 output into another team's Deliverable 2. Does this imply Deliverable 2 must accept any standard skeleton format, or will a common keypoint schema be specified before TD1?

    Answer: IARPA and the T&E team, taking into account the technical approaches of the selected Performers and evaluating the costs and benefits of doing so, may require a Team's Technical Deliverable 2 to conform to a common interface standard.

218. MOVES: Please clarify whether IARPA funds may be used to support data collection by performers, beyond the T&E-provided cohort, including:

    a. Recordings on additional subjects.

    b. Associated costs — Human Subject Research / IRB review, equipment, subject compensation, data-handling infrastructure, clinical labeling.

    c. If permitted, are there scope, volume, or HSR-oversight requirements (e.g., performer-institution IRB versus program-level review)?

       Answer: Performers may collect data, following all federal HSR guidelines and requirements. Typically, this can be satisfied by performer-institution IRB with a federal wide assurance.

219. MOVES: Regarding the Algorithm 2 performance target of 80% diagnostic accuracy on chosen disorders, using a balanced test set and a user-settable threshold selected post hoc, please clarify the following:

    a. Is the 80% target applied independently to each chosen disorder as a one-vs-all binary evaluation, or as a single multi-class evaluation across the chosen disorders and healthy controls?

    b. If the target is evaluated as a multi-class task, what is the operational definition of "accuracy"? For example, does it refer to overall argmax top-1 accuracy, balanced accuracy meaning mean per-class recall, per-class recall, or another metric?

       Answer: The accuracy metric is overall top-1 accuracy.

    c. If the target is evaluated as a one-vs-all task, what is the composition of the negative class? Does it include healthy controls only, healthy controls plus other chosen disorders, or another group?

       Answer: N/A

    d. How is the balanced test set balanced? For example, is it balanced 50/50 between target and all in the one-vs-all case, equally across all classes in the multi-class case, or according to another principle?

       Answer: Best efforts will be made to have rough balance between people with movement disorders and no known movement disorders. Within those with movement disorders, best efforts will be made to have a roughly equal distribution of disorders.

    e. Is the 80% metric computed at the clip level, with each 15-second output clip treated as an independent unit, or at the subject level, with clips aggregated into a per-subject decision before scoring?

       Answer: It is computed at the subject level.

220. ARCADE: The program description indicates that final architectures should remain adaptable for customers requiring self-hosted, closed-network deployments, and that Performers are encouraged to leverage state-of-the-art LLMs during development.

    Could you describe the environment in which the T&E team intends to execute performer-delivered containers (e.g., commercial cloud, AWS GovCloud / Azure Government, on-premise GPU cluster, air-gapped enclave)? Understanding the compute, network, and dependency constraints would let us right-size our containerization and dependency strategy.

    Answer: Compute environments will be discussed at kick-off.

    For the "adaptable to self-hosted, closed-network deployment" requirement, would environments such as AWS GovCloud or Azure Government satisfy this, or does the government anticipate fully air-gapped on-premise targets? This would

affect whether we can rely on commercial cloud-managed services in any portion of the reference architecture.

Answer: This is currently unknown. Efforts should concentrate on the best possible solution.

Are there compute or budget constraints that we should anticipate for inference at T&E time (e.g., max GPU count, memory budget, etc.)

Answer: Currently unknown. Additional information will be provided at kick-off.

221. MOVES: If after May 27th the programs are still active, will advance notification be given before the program becomes inactive? If yes, will there be an expected window (ex. 72 hours in advance) for notification? (ETA Program Announcement - Section 4.A.1)

Answer: No, if there is no Amendment re-labeling the ETA as inactive, assume it is still active. This is applicable to all ETA Programs.

222. MOVES: Are non-US based foreign employees of a US company eligible to perform on MOVES? For example, the subcontracting company is a US company but the specific personnel to perform resides outside the US and are not US citizens. (ETA Program Announcement - Section 3.A)

Answer: The Section cited in the question states:

Foreign entities and individuals may participate as subcontractors or employees of a U.S. entity; however, all foreign participation must comply with any necessary Non-Disclosure Agreements, Security Regulations, Export Control Laws, and other governing statutes applicable under the circumstances. Proposers are responsible for ensuring that the efforts of foreign participants do not either directly or indirectly compromise the laws of the United States, nor its security interests.

223. MOVES: Does this program require NIST 800-171 / CMMC certification as part of the Research Security Program (RSP)? (ETA Program Announcement - Section 5.C)

Answer: We do not anticipate requiring NIST 800-171 or CMMC certification for MOVES performers

224. Is the MOVES program subject to export control? (ETA Program Announcement - Section 6.B.6)

Answer: While we do not anticipate export control, a final determination for each solution will be evaluated independently at the time of delivery.

225. MOVES: Are there hardware specifications that the solution is expected to perform under? (ETA Programs - MOVES - "Program Technical Deliverables and Performance Metrics")

Answer: No.

226. MOVES: Will the Government provide a cloud-based development environment for performer teams to use on the program? (ETA Programs - MOVES - "Government Furnished Equipment/Information")

Answer: No.

227. COSMIC: Is the program subject to export control? (ETA Program Announcement - Section 6.B.6)

Answer: While we do not anticipate export control, a final determination for each solution will be evaluated independently at the time of delivery.

228. COSMIC: Does this program require NIST 800-171 / CMMC certification as part of the Research Security Program (RSP)? (ETA Program Announcement - Section 5.C)

Answer: We do not anticipate requiring NIST 800-171 or CMMC certification for COSMIC performers.

229. COSMIC: Will the Government provide a cloud-based development environment for performer teams to use on the program? (ETA Programs - COSMIC - "Test & Evaluation")

Answer: Yes, some cloud computing resources will be provided by the Government. Specific hardware to train, develop, and test hardware is an allowable ODC to be included in the Cost Volume.

230. COSMIC: One objective is the ability to identify ephemeral events vs persistent change but development of change detection methods is marked as out of scope. Could you clarify where the identification of ephemeral events/persistent change would be in use? Would it function like an auxiliary classification to T&E provided change detection algorithm(s)? (ETA Programs - COSMIC - "Objectives")

Answer: The identification of ephemeral events versus persistent change is a potential application of the PPD. PPDs will be tested to assess their ability to discriminate between these changes using benchmark algorithms.

231. LOCUS: Is the program subject to export control? (ETA Program Announcement - Section 6.B.6)

Answer: While we do not anticipate export control, final determination for each solution will be evaluated independently at the time of delivery.

232. LOCUS: Does this program require NIST 800-171 / CMMC certification as part of the Research Security Program (RSP)? (ETA Program Announcement - Section 5.C)

Answer: We do not anticipate requiring NIST 800-171 or CMMC certification for LOCUS performers.

233. LOCUS: Would other non-ground sensor data such as satellite or aerial imagery be provided as "reference data"? (ETA Programs - LOCUS - "Program Specific Terms/Definitions")

Answer: Information about reference data will be provided at kick-off.

234. LOCUS: What kind of reference data would be provided for indoor scenarios to aid geolocation of that query? (ETA Programs - LOCUS - "Program Specific Terms/Definitions")

Answer: Information about reference data will be provided at kick-off.

235. LOCUS: Will the Government provide a cloud-based development environment for performer teams to use on the program? (ETA Programs - LOCUS - "Government Furnished Equipment/Information")

Answer: No. Proposers are expected to perform their own integration into the test and evaluation cloud platform. ODCs for specialized hardware for core development are available.

236. LOCUS: Do you have plans to score geolocation explainability? (ETA Programs - LOCUS - "Program Metrics")

Answer: The explainability assessment will be discussed during kick-off.

237. LOCUS: Will evaluation consider a geolocation successful if and only if it can be explained? In other words, if a geolocation is accurate but there is no explanation offered, or it is nonsense, will it be considered unsuccessful or partially successful? (ETA Programs - LOCUS - "Program Scope and Limitations")

Answer: Geolocation explainability will be assessed independently from geolocation accuracy, but a holistic solution includes both.

238. LOCUS: Are there hardware specifications that the solution is expected to perform under? (ETA Programs - LOCUS - "Software Deliverable Formatting")

Answer: Additional information will be provided at kick-off and discussed throughout the program.

239. LOCUS: Does this system have to run isolated on the provided hardware or can it connect to external API and datasets?

Answer: The system is expected to operate at full performance in an air gapped environment.

240. LOCUS: How far ahead of time will we receive the reference data?

Answer: Information about reference data will be provided at kick-off.

241. LOCUS: Can you provide additional information about what the reference data will include?

Answer: Information about reference data will be provided at kick-off.

242. LOCUS: Are there any hardware specifications for what kind of hardware will the T&E team run our models on? Do we have any constraints? Will they be able to provide GPUs if our system needs on GPUs

Answer: Additional information will be provided at kick-off and discussed throughout the program.

243. LOCUS: Does explainability have to be plain language or can something like visual explanation grad-CAM work?

Answer: Plain language is preferred, but other explainability methods are also desirable.

244. LOCUS: Do each image and audio models have to provide plain explainability or is it OK as long as the final combined system provides plain language explainability?

Answer: Plain language is preferred, but other explainability methods are also desirable.

245. LOCUS: Are the audio channel stereo or mono channel for the dataset?

Answer: Both

246. LOCUS: Does the system have to run autonomously right from the start, or can we have a model that requires human-in-the-loop initially, as long as it is autonomous by the end of the program?

Answer: It needs to be autonomous for the test events.

247. LOCUS: Can we include a subcomponent that is a commercial system accessible via a license only, rather than code transfer to the govt? Licenses will be provided to the government throughout the program, but they may negotiate a contract after the program to continue licensing.

Answer: Reference answer to Question # 109B above.

248. DECIPHER: Regarding T&E, is it anticipated that span identification will be expected from candidate systems? If so, will non-contiguous morphology be considered out of scope? (e.g., infixes, circumfixes, transfixes)

Answer: Span identification is in scope. Non-contiguous morphology is not out of scope.

249. DECIPHER: Will T&E score partially overlapped span predictions on a scale or as a correct/incorrect binary?

Answer: Additional information will be provided during the program.

250. DECIPHER: Will pre-processing be applied to test set items and if so, at what stage? Will performers have access to the raw and processed test set items?

Answer: Additional information may be provided during the program.

251. DECIPHER: Should we assume that generated definitions will be in English regardless of source language, or should the generated definitions be in the respective source languages?

Answer: Yes, generated definitions will be in English regardless of source language.

252. DECIPHER: Will duplicate examples of coded language be removed from the test set or should we expect duplicates?

Answer: Additional information may be provided during the program.

253. COSMIC - The announcement says that solutions "should rapidly be able to ingest new imagery, data, and data types..." At what stage should we expect to incorporate new types of data? Will we have to incorporate previously unknown data types during evaluation, or will we know the data types prior to submitting the system?

Answer: New data sources will be introduced at the PPD creation and adaptation stage. It should be anticipated that the Government will test performance against the ingestion of unknown data types.

254. COSMIC - In the proposer's day slides, an example was shown of how users and agents might interact with the PPD source. One example makes a request for "High-Res Tanks." Does our solution need to be able to classify and detect objects in the data, or does "tank" refer to information available in the metadata or other data sources?

Answer: This was an example of a potential application of the PPD

255. COSMIC - When converting between modalities, do we want to generate an output that is realistic or idealized? For instance, if we are converting from SAR to EO and we know there is light cloud cover, should we generate an image obscured by clouds or an ideal image as if there were no clouds?

Answer: Realistic.

256. LOCUS: What is the expected use of Reference Data or representative materials supplied by T&E before test events? Is Reference Data intended for pre-test preparation and interoperability only, or may performer systems access it at test time?

Answer: Reference data will be discussed at program kick-off.

257. LOCUS: Is multi-container orchestration (Docker Compose, Airflow DAG) permitted, or single-container only? What is the maximum on-disk size for the container plus bundled reference data?

Answer: Multi-container orchestration is permissible. Details of the computational environment will be discussed at kick-off.

258. LOCUS: Will a reference container or integration test harness be provided to performers before Test Event 1, and at what point in the Pre-program Phase?

Answer: Exemplar data and details of the computational environment will be discussed at kick-off.

259. LOCUS: What runtime environment should we assume for T&E? Specifically: (a) CPU/GPU/memory/storage available per performer, (b) whether the container is air-gapped or has network access to external APIs and commercial AI services at inference time, (c) if air-gapped, the size budget for bundled reference data and whether it can be staged separately from the container image, and (d) whether the 5-minute Time to Prediction is measured inclusive of reference-data lookup and any allowed network calls.

Answer: Details of the computational environment will be discussed at kick-off.

260. LOCUS: Is the 5-minute Time to Prediction per query regardless of media length, or scaled by media duration? Does it include media ingestion, reference-data lookup and explanation generation, or only inference? The footnote says it applies to a continent-scale search area; what is the time target for smaller search areas?

Answer: Details of metric computation will be provided at program kick-off.

261. LOCUS: What output format will the LOCUS API standard require: single point, ranked candidates with confidences, probability distribution, polygon, or selectable? What confidence representation is expected, and how does it interact with CE90?

Answer: Details of output format will be discussed at kick-off and during the program.

262. LOCUS: Do we need to create a user interface for analysts, or will T&E interact with the prototype only through a containerized API or command-line interface?

Answer: User interfaces are not necessary during initial program activities.

263. LOCUS: How will Geolocation Explainability be evaluated? Will it be scored quantitatively, reviewed qualitatively, or treated as a required output without a separate score?

Answer: The explainability assessment will be discussed during kick-off.

264. LOCUS: What minimum capability should we deliver for the first software delivery and first test event? Should the Month 2/TE1 prototype already be a complete end-to-end multimodal system?

Answer: Specific program deliverables will be discussed during program kick-off.

265. LOCUS: How are capabilities expected to mature across the five test events? For example, may we phase in stronger audio, video, indoor, outdoor, reference-data, and explainability capabilities over time?

Answer: We intend to increase the difficulty of test events as time goes on.

266. LOCUS: The announcement says T&E may supply reference data "with certain qualities or characteristics that will be representative of queries of interest." Will reference data for a given Test Event be delivered far enough in advance to allow pre-processing into performer-internal representation at scale, and will it ever be delivered in the same package as the queries themselves?

     Answer: Information about reference data will be provided at kick-off.

267. COSMIC: Will performers be expected to address all methods on page 2 of the draft announcement by Month 3. For the first 3-month software deliverable, it seems early to expect a PPD delivery for testing. For T&E, would you expect a containerized package to test the interfaces with the T&E tools to be sufficient? Or will the program add methods as the program progresses? If the latter, does IARPA have a prioritization of capabilities they want to see/test?

     Answer: Expectations for the first challenge event delivery will be provided at program kick-off. The program anticipates expanding testing to determine increasing functionality and robustness of the PPD throughout the program.

268. COSMIC: What portion of a performer's budget do you consider acceptable for dataset acquisition?

     Answer: Cost estimates will be evaluated on a case-by-case basis.

269. COSMIC: To avoid duplication of dataset expenses among performers, might IARPA consider data sharing among performers?

     Answer: All data acquired by Performers is expected to be deliverable to the Government and may be shared with performers.

270. COSMIC: For the interim and final demos, will the testing be focused on a specific manageable global subset, perhaps on a small geographic region (Like Jacksonville Florida or Washington DC) and with a subset of potential geospatial data types?

     Answer: Testing will be on a set of diverse geographic locations.

271. COSMIC: There seems to be separate technical development areas. For example, the design, development and construction of the PPD, the data harmonization. At what stage for the incremental deliveries will each of these need to be included in the deliveries and at what stage should these be integrated into a delivery?

     Answer: Specifics of expectations for deliveries for challenge events will be specified at program kick-off. We do anticipate deliveries of PPDs early in the program and expect them to continually advance throughout the program duration.

272. COSMIC: Do you have an example use cases that you are most interested in seeing the performers upon which to focus?

     Answer: Example use cases are available in the proposer’s day slides and topic description.

273. COSMIC: The draft solicitation mentions that data will be provided. Can you give us an idea of what data types and volume expect to be included in this delivery?

Answer: Exemplar data from a variety of commercial remote sensing sources and other geolocated data will be provided during the life of the program. This data should not be of volume sufficient for training.

274. COSMIC: Please clarify the due dates for the COSMIC white papers. From the RFP it sounds like it will be determined after the Q&A responses and when a sufficient number of white papers are received.

Answer: Window to submit White Papers is open now and is guaranteed to stay open until 27 May 2026. This is applicable to all ETAs.

275. COSMIC: Will some airborne and satellite data be provided by IARPA to complete the data required for COSMIC development and evaluation?

Answer: Exemplar data from a variety of commercial remote sensing sources and other geolocated data will be provided during the life of the program. This data should not be of volume sufficient for training.

276. ARCADE: We are working to identify a prime for one of our submissions. Our intent was to submit two white papers until we were informed this week this was not a viable option. If we proceed to submit two white papers as a prime, what are the potential outcomes?

Answer: Submit one as a Prime until you hear from the AO. If determined not viable and assuming the ETA is still active, submit the second one.

277. ARCADE: Are there any follow-on security considerations for subsequent transition of capability to the IC customer…. including US citizenship of prime / sub PIs on the white papers / proposals?

Answer: This is currently unknown.

278. DECIPHER: Is internet access permitted in the T&E environment, and if not, are closed-PI commercial LLMs effectively excluded from final delivered systems?

Answer: While some internet access may be available in the T&E environment, we will also be testing performance on stand-alone systems. This may preclude the incorporation of closed-API models

279. DECIPHER: What is the approximate size, document-length distribution, and number of languages in the sequestered test set per evaluation at months 3, 6, 9, and 12?

Answer: Additional information may be provided during the program.

280. DECIPHER: For Definition Generation, is ground truth a single reference per term or multiple acceptable definitions, and how is semantic equivalence judged (LLM-as-judge, human, or both)?

Answer: The test and evaluation approach will be discussed during the course of the program.

281. DECIPHER: Per Q42, will low-resource languages appear in scored T&E sets, or is “applicability” demonstrated only via performer-provided evidence?

     Answer: The program may test on low resource languages.

282. DECIPHER: For concept drift detection, will the test corpora themselves span a sufficient time range to observe drift in-corpus, or are performers expected to maintain external longitudinal reference data (e.g., historical snapshots of a term's usage) to detect drift relative to a prior baseline?

     Answer: We are looking for creative and feasible approaches to detect and measure concept drift.

283. DECIPHER: Will timestamp metadata be provided at the document level, corpus level, or both, and will timestamps reflect authentic publication/communication dates or synthetic dates assigned for evaluation? This affects both social/contextual annotation and concept drift detection.

     Answer: Timestamp metadata may or may not be included during testing.

284. DECIPHER: The announcement specifies turnkey operation with no human-in-the-loop. Within that constraint, is automated test-time adaptation by the delivered system permitted (e.g., in-container fine-tuning, few-shot in-context learning, or unsupervised adaptation on the test corpus using the provided domain description), or must model weights remain frozen at delivery?

     Answer: Automated test time adaptation would be acceptable; however, we may require frozen weights to facilitate transition.

285. ARCADE: Could the Government provide a rough estimate of the volume and specific sub-categories of the initial component library dataset?

     Answer: The library will include a wide variety of commercially available components, including passive, active, electromechanical, RF, power-management, sensors, connectors, microcontroller families, and some FPGAs.

286. ARCADE: Could the Government provide specifications on the hardware constraints (e.g., GPU VRAM limits, node count) for the T&E deployment environment, and how exactly the 'baseline system' is defined.

     Answer: Compute environment details will be provided at kick-off.

287. ARCADE: Could the Government clarify if there are specific Country of Origin restrictions, software supply chain security requirements, or prohibited foreign-developed model lists that apply to the open-source or open-weight foundational models packaged within the Performer's containerized deliverable?

Answer:

PROHIBITION ON CONTRACTING FOR CERTAIN TELECOMMUNICATIONS AND VIDEO SURVEILLANCE SERVICES OR EQUIPMENT

(a) Definitions.

Backhaul means intermediate links between the core network, or backbone network, and the small subnetworks at the edge of the network (e.g., connecting cell phones/towers to the core telephone network). Backhaul can be wireless (e.g., microwave) or wired (e.g., fiber optic, coaxial cable, Ethernet).

Covered foreign country means The People's Republic of China.

Covered telecommunications equipment or services means—

(1) Telecommunications equipment produced by Huawei Technologies Company or ZTE Corporation (or any subsidiary or affiliate of such entities);

(2) For the purpose of public safety, security of Government facilities, physical security surveillance of critical infrastructure, and other national security purposes, video surveillance and telecommunications equipment produced by Hytera Communications Corporation, Hangzhou Hikvision Digital Technology Company, or Dahua Technology Company (or any subsidiary or affiliate of such entities);

(3) Telecommunications or video surveillance services provided by such entities or using such equipment; or

(4) Telecommunications or video surveillance equipment or services produced or provided by an entity that the Secretary of Defense, in consultation with the Director of National Intelligence or the Director of the Federal Bureau of Investigation, reasonably believes to be an entity owned or controlled by, or otherwise connected to, the government of a covered foreign country.

Critical technology means—

(1) Defense articles or defense services included on the United States Munitions List set forth in the International Traffic in Arms Regulations under subchapter M of chapter I of title 22, Code of Federal Regulations;

(2) Items included on the Commerce Control List set forth in Supplement No. 1 to part 774 of the Export Administration Regulations under subchapter C of chapter VII of title 15, Code of Federal Regulations, and controlled—

(i) Pursuant to multilateral regimes, including for reasons relating to national security, chemical and biological weapons proliferation, nuclear nonproliferation, or missile technology; or

(ii) For reasons relating to regional stability or surreptitious listening;

(3) Specially designed and prepared nuclear equipment, parts and components, materials, software, and technology covered by part 810 of title 10, Code of Federal Regulations (relating to assistance to foreign atomic energy activities);

(4) Nuclear facilities, equipment, and material covered by part 110 of title 10, Code of Federal Regulations (relating to export and import of nuclear equipment and material);

(5) Select agents and toxins covered by part 331 of title 7, Code of Federal Regulations, part 121 of title 9 of such Code, or part 73 of title 42 of such Code; or

(6) Emerging and foundational technologies controlled pursuant to Section 1758 of the Export Control Reform Act of 2018 (50 U.S.C. 4817).

Interconnection arrangements means arrangements governing the physical connection of two or more networks to allow the use of another's network to hand off traffic where it is ultimately delivered (e.g., connection of a customer of telephone provider A to a customer of telephone company B) or sharing data and other information resources.

Reasonable inquiry means an inquiry designed to uncover any information in the entity's possession about the identity of the producer or provider of covered telecommunications equipment or services used by the entity that excludes the need to include an internal or third-party audit.

Roaming means cellular communications services (e.g., voice, video, data) received from a visited network when unable to connect to the facilities of the home network either because signal coverage is too weak or because traffic is too high.

Substantial or essential component means any component necessary for the proper function or performance of a piece of equipment, system, or service.

(b) Prohibition.

(1) Section 889(a)(1)(A) of the John S. McCain National Defense Authorization Act for Fiscal Year (FY) 2019 (Pub. L. 115-232) prohibits the head of an

executive agency on or after August 13, 2019, from procuring or obtaining, or extending or renewing a contract to procure or obtain, any equipment, system, or service that uses covered telecommunications equipment or services as a substantial or essential component of any system, or as critical technology as part of any system. The Performer is prohibited from providing to the Government any equipment, system, or service that uses covered telecommunications equipment or services as a substantial or essential component of any system, or as critical technology as part of any system, unless an exception at Paragraph (c) of this clause applies or the covered telecommunication equipment or services are covered by a waiver.

(2) Section 889(a)(1)(B) of the John S. McCain National Defense Authorization Act for FY 2019 (Pub. L. 115-232) prohibits the head of an executive agency on or after August 13, 2020, from entering into a contract, or extending or renewing a contract, with an entity that uses any equipment, system, or service that uses covered telecommunications equipment or services as a substantial or essential component of any system, or as critical technology as part of any system, unless an exception at Paragraph (c) of this clause applies or the covered telecommunication equipment or services are covered by a waiver. This prohibition applies to the use of covered telecommunications equipment or services, regardless of whether that use is in performance of work under a Federal contract.

(c) Exceptions. This clause does not prohibit contractors from providing—

(1) A service that connects to the facilities of a third-party, such as backhaul, roaming, or interconnection arrangements; or

(2) Telecommunications equipment that cannot route or redirect user data traffic or permit visibility into any user data or packets that such equipment transmits or otherwise handles.

(d) Reporting requirement.

(1) In the event the Performer identifies covered telecommunications equipment or services used as a substantial or essential component of any system, or as critical technology as part of any system, during ETA performance, or the Performer is notified of such by a subcontractor at any tier or by any other source, the Performer shall report the information in Paragraph (d)(2) of this clause to the AO.

(2) The Performer shall report the following information pursuant to Paragraph (d)(1) of this clause:

(i) Within one business day from the date of such identification or notification: The Agreement number; the order number(s), if applicable; supplier name; supplier unique entity identifier (if known); supplier Commercial and Government Entity (CAGE) code (if known); brand; model number (original equipment manufacturer number, manufacturer part number, or wholesaler number); item description; and any readily available information about mitigation actions undertaken or recommended.

(ii) Within 10 business days of submitting the information in Paragraph (d)(2)(i) of this clause: Any further available information about mitigation actions undertaken or recommended. In addition, the Performer shall describe the efforts it undertook to prevent use or submission of covered telecommunications equipment or services, and any additional efforts that will be incorporated to prevent future use or submission of covered telecommunications equipment or services.

(e) Subcontracts. The Performer shall insert the substance of this clause, including this Paragraph (e) and excluding Paragraph (b)(2), in all subcontracts and other contractual instruments, including subcontracts for the acquisition of commercial items.

288. COSMIC: Are change detection algorithms the only downstream algorithms used during test events, or will other downstream algorithms (e.g., agentic AI, object detection, classification, retrieval) also be exercised against the delivered PPD?

Answer: Delivered PPDs will be tested against a variety of downstream algorithms, including change detection.

289. COSMIC: The solicitation lists "answer fidelity" as a metric. What downstream algorithms or workflows will be used to evaluate answer fidelity, and does the delivered system need to expose a chat or query interface for these agentic AI systems?

Answer: A series of benchmark questions requiring persistent surveillance will be curated to test answer fidelity. The expectation is that the harmonized PPD will allow interrogation by an Agentic AI system.

290. COSMIC: Is ground-level imagery (e.g., from ground-based cameras, dashcams, body cameras, or social-media photos) expected as input to the system, and is such imagery expected as an output or queryable layer of the PPD?

Answer: Geolocated ground level imagery may be expected to be incorporated into the PPD.

291. COSMIC: What compute is available in the sequestered T&E environment (GPU type and count, VRAM, RAM, storage, runtime budget per query or per evaluation)?

Answer: Details about computational requirements for T&E will be provided at program kick-off.

292. COSMIC: Is internet access permitted in the T&E environment, and if not, are closed-API commercial LLMs or geospatial services effectively excluded from final delivered systems?

Answer: While some internet access may be available in the T&E environment, we will also be testing performance on stand-alone systems. This may preclude the incorporation of closed-API models.

293. DECIPHER: The TA2 specification calls for definitions following lexicographic best practices while simultaneously requiring the system to work on small collections of content containing limited attestations of a term. Established lexicographic methodology typically requires a minimum corpus density per sense before a definition is considered well-formed. When a specialized term in the test set has attestations below that threshold, will T&E credit a context-grounded best-effort definition that the system marks with appropriate confidence, or will only definitions meeting full lexicographic standards score positively?

Answer: Further information about test and evaluation will be shared during the program. Achieving full lexicographic standards may not be feasible or desirable in the DECIPHER program.

294. DECIPHER: The Detection task notes that the threshold between specialized and non-specialized language may differ depending on the end user's knowledge of a domain. Will T&E evaluation be conducted against a single fixed domain-knowledge profile, or across multiple profiles representing different baseline expertise levels? If single, how is that profile characterized; if multiple, will Performer systems be scored on per-profile precision/recall or on aggregate?

Answer: Further information about test and evaluation will be shared during the program.

295. DECIPHER: The announcement constrains the test set to high- and medium-resource languages. For Performer systems where multilingual coverage varies sharply by language family (Indo-European versus Sino-Tibetan versus Afro-Asiatic, for instance), will the language-family composition of the test set be disclosed at program kick-off, or only at test event?

Answer: The language family composition of the test set may not be shared, but information about the test and evaluation approach including specific languages of interest will be discussed during the program.

296. DECIPHER: The TA1+TA2 interface is described as user-friendly and is intended for use by IC analysts. Section 508 and WCAG 2.1 AA compliance are statutory requirements for IC analyst-facing software. Will compliance be required at the M11 deployable software delivery, or is functional-prototype-only acceptable with accessibility remediation deferred to transition? If required at delivery, will the standard be Section 508 Revised (ICT) baseline, WCAG 2.2 AA, or a specific IC-internal accessibility standard?

Answer: We are looking for a prototype; additional software compliance requirements are not in scope at this phase.

297. DECIPHER: The Program Overview describes detecting concept drift by monitoring use of terminology as it evolves over time, but neither TA1 nor TA2 makes concept drift a discrete task or a metric. Is concept drift detection treated as an emergent capability to be qualitatively assessed, a separate quantitative metric the Government may add, or out of scope for the M3, M6, M9, and M12 evaluation cycle?

    Answer: We are looking for creative and feasible approaches to detect and measure concept drift.

298. LOCUS: Are we allowed to use reference datasets such as Satellite imagery during the inference/ testing process of our system? I understand we provide you with a Docker container of our code. But can that also include the reference datasets of satellite imagery?

    Answer: Yes.

299. MOVES: What is the Government's definition of "unclothed" as it pertains to Technical Deliverable #1 software simulation?

    Answer: The unclothed full-body animation is intended to illustrate the motion of the body unobstructed by clothing, for clinician review in the context of diagnosing movement disorders. It should not be photorealistic or include unnecessary details. Identifiability of the real individual from the unclothed body animation is not necessary.

300. MOVES: Does the Government require the software simulation to be photorealistic, or is rendering via a computer graphics engine preferable?

    Answer: The unclothed full-body animation is intended to illustrate the motion of the body unobstructed by clothing, for clinician review in the context of diagnosing movement disorders. It should not be photorealistic or include unnecessary details. Identifiability of the real individual from the unclothed body animation is not necessary.