

Biometric Recognition and Identification at Altitude and Range (BRIAR) Program					
IARPA-BAA-20-04 - AMD #4					
IARPA will not be accepting any further questions on BAA 20-04					
Industry Questions and Government Answers - ROUND 2 - January 28, 2021					
Question number	Industry Questions	Paragraph & Page # (BAA-A3)	Government Responses	Change to BAA (Yes/No)	Paragraph & Page # (BAA-A4)
9	Is there an address to which a 1-page white paper may be submitted for feedback on the ideas/goals?	N/A	IARPA is not accepting any white papers as part of the BRIAR BAA.	No	N/A
10	The proposal instructions include the RDMP in both Section L (which counts against the 40 page limit) and Attachment 11 (which does not count against the 40 page limit). Is the full 2-3 page RDMP required in both of these section, or can we omit Section L and just include the RDMP as Attachment 11?	4.B.1.c, p. 37-38	The proposal requires a "Detailed Management Plan" and a "Research Data Management Plan", (RDMP). These are two separate things. The Detailed Management Plan is included in the 40 page proposal limit and the RDMP is a proposal attachment – not included in the page count for the 40 page limit.	No	N/A
11	We currently have a contract with ODN/IARPA. Are we prohibited to participate in this effort?	3.A.1, p. 30-31	Submission of a proposal is a business decision. Per Section 3.A.1 Organizational Conflicts of Interest (OCI) - In accordance with FAR 9.5, Offerors are required to identify and disclose all facts relevant to potential OCIs involving the Offeror's organization and any proposed team member (subawardee, consultant). Under this Section, the Offeror is responsible for providing this disclosure with each proposal submitted pursuant to the BAA.	No	N/A
12	Can USG provide further clarification on 'identity and sighting instance labels' specified as pieces of information needed for annotation, specifically the 'sighting instance' portion?	1.D.1.2, p. 15	An identity label is associated with a specific individual; a sighting instance label is a unique label for each person within each video clip across the dataset. Note that a long video could have multiple sighting instances for the same individual if temporally separated by a gap in viewing.	Yes	1.D.1.2, p. 15
13	Can USG provide guidance on annotation parameters or methodologies to ensure interoperability (e.g. bounding box attributes)?	1.D.1.2, p. 15	Annotation methodologies should be consistent with current research best practices to facilitate supervised machine learning algorithm development. Additional guidance on annotation standards and methodologies are under development by T&E and will be provided at program kickoff.	Yes	1.D.1.2, p. 15
14	Are former IARPA datasets (IJB-A, IJB-B, IJB-C, IJB-S) eligible for use?	1.D.1.3, p. 15	Yes, at this time IARPA will allow the use of Janus datasets for developmental purposes, but this decision is subject to change as the BRIAR T&E objectives mature. Additional direction will be provided at program kickoff. IARPA will facilitate access by Performers to IJB-S and IJB-MDF upon program start.	No	N/A
15	In Table 2, it appears that for each performer dataset, each team needs to collect a dataset of 200 individuals for the Probe dataset and another 200 distractor individuals for the Gallery dataset. However, in the text on p. 14 it says that each performer dataset must consist of 200 unique consenting individuals. Is the Gallery dataset made up of another 200 unique consenting individuals, thus making each performer data collection require 400 unique individuals?	1.D, p. 10 1.D.1.2, p. 14	No, 200 unique subjects <b>total</b> are required for each BRC dataset and will appear in both the Probe and Gallery data. BTS datasets will contain additional distractor subjects in their galleries.	Yes	1.D, p. 10 (footnote)
16	For BRC datasets, should the Gallery data be collected in the same way as Probe data (i.e., at least 4 sensors per individual, 2 sets of clothing per individual, etc.)?	1.D.1.2, p. 14	Imagery collected for probe and gallery are separate data collection events and governed by different conditions. Gallery enrollment collections should follow industry standards and research best practices to facilitate standardized, multi-view still and video imagery of subjects to support face and whole-body biometric matching (static and gait motion).	Yes	1.D.1.2, p. 14-15
17	For BRC datasets, does the Gallery dataset need to be annotated like the Probe dataset?	1.D.1.2, p. 15	Yes, per 1.D.1.2, "Performers will be responsible for annotating all BRC data."	No	N/A
18	At the ~50TB/dataset mentioned in the BAA, storage for program datasets would total 1PB. Storing this volume of data will be more expensive if performers need to backup the data within each team. Can performers assume that program data will have persistent availability, e.g. will the BRS1 data made available in month 1 continue to be available via Globus throughout the program?	1.D, p. 10	Yes, Performers will have persistent access to all relevant program data through the BRIAR T&E data repository.	No	N/A
19	For the purpose of planning the complexity of our data collection, to what extent do BRCs need to cover all combinations of subjects, sensors, and viewing conditions (distance and pitch)?	1.D.1.2, p. 13-14	The threshold (i.e., minimum) requirement for BRCs is that they support their current Phase's objectives and goals, with the objective requirement (i.e., desired goal) being that they support all Phases' objectives and goals. For example, the threshold requirement for BRC1 is to support Phase 1 needs, such as moderate range and controlled indoor enrollment, but it is desirable to also collect severe range imagery and partially constrained enrollment imagery.  The representation frequency of imagery conditions within a given BRC is flexible, requiring only that each basic condition is present. Performers have flexibility in designing each BRC and its resulting imagery to best support their research approaches, address areas that may need additional data to overcome limitations, and to meet overall program goals. Note, that minimum requirements for BRCs listed in Section 1.D.2 (p. 14) are still valid and need to be met.	Yes	1.D.1.2, p. 15
20	The BAA specifically refers to "quad-copters." We understand that a list of government preferred UAV platforms will be provided during the program, but does this wording mean that other rotorcraft UAVs (hex-copters) are not of interest?	1.D.1.1, p. 12 1.D.1.2, p. 14	Other rotorcraft UAV, such as hexcopters and helicopters, are relevant and in scope.	Yes	1.D.1.1, p. 12 1.D.1.2, p. 14
21	The BAA states that the subject counts in a probe are a single subject for Phase 1 (1.A.2.1), up to three for Phase 2 (1.A.2.2), and "unconstrained" for Phase 3 (1.A.2.3). Section 1.D.1.2 states that BRC collections should have up to three persons per scene. Is there an expectation that BRCs should include data capable of creating probes with subject counts up to the applicable phase constraint and therefore BRC4 (Phase 3) should include scenes with more than three people?	1.A.2.1, p. 6 1.A.2.2, p. 7 1.A.2.3, p. 8 1.D.1.2, p. 14	See answer for Question #19.	Yes	1.D.1.2, p. 15

22	<p>This question concerns the "Subjects" columns in Table 2, and the number of subjects to recruit for each BRC.</p> <p>a. The BAA stipulates (e.g. page 14) that BRCs should have a minimum of 200 unique subjects, clarified by footnote 16 (page 14) that each BRC should include at least 200 new individuals.</p> <p>b. Table 2 states that (for example) BRC1 should have 800 "Subjects (Probes)" and 800 "Subjects (Gallery)", which footnote 10 defines as "Minimum number of unique individuals in the probe set with mated entries in the dataset gallery".</p> <p>We had been assuming that each individual recruited for a BRC (a "subject") would (1) map to a single entry in the BRC gallery with associated enrollment data, and (2) also generate multiple probe entries generated from data collected during the event. It seems either (a) BRC1 in fact would require 800 recruits for an 800-subject gallery, not 200, or (b) subjects appear multiple times in the gallery, or (c) the BRC gallery should contain individuals not recruited as part of the BRC; the origin of these individuals is not discussed, or (d) the "800 Subjects (Gallery)" entry is incorrect and should be 200.</p>	1.D, p. 11	BRC1 in Table 2 is the aggregation of all BRC1 datasets collected by the Performers. Per Footnote #14, Table 2 numbers assume four Performers in Phase 1. BRC1 for each Performer will contain a minimum of 200 subjects.	Yes	1.D, p. 11 (footnote)
23	Can you clarify any expectations about how BRC imaging events should be distributed between probes and the gallery? Footnote 17 (page 14) says that each of the BRC subjects should be imaged by at least four different sensors, but does not indicate whether these should be simultaneous or independent, or how these imaging events should be allocated between the gallery and the probe set.	1.D.1.2, p. 14	See answer to Question #16.  Within BRCs, probe imagery of subjects from different sensors are not required to be simultaneous. Enrollment video for gait under controlled and partially controlled conditions shall be captured simultaneously at multiple camera angle views.	Yes	1.D.1.1, p. 12-13 1.D.1.2, p. 14-15
24	For a given clip in which a subject is visible, is it expected that the clip also contain video where the subject is not visible? The question arises when addressing the logistics of collecting BAA-compliant data containing "video of varying durations (e.g. seconds up to minutes)" (page 14). For a goal of (say) 10 instances of clips of 12 seconds each, one way is to collect a single 120 second clip of the subject and slice it into 10 segments. Another way is to separately script 10 individual short takes in which the subject is off-scene, walks on-scene for 12 seconds, then walks off-scene. The former is much easier to obtain but lacks the transitions of the latter.	1.D.1.2, p. 14	Video clips are not required to contain video without the subject. Both approaches to generating multiple video sample instances are acceptable. Government Research Sets may include either or both types of video clips. Performers should plan their data collection approaches to meet their research approach needs.	No	N/A
25	Privacy Plan - Is there an expectation that we anonymize the identities of the volunteers in researcher-collected datasets (BRC)? If so, will that extend to the online presence of those volunteers on various platforms?	1.D.3, p. 17	Yes, anonymization of subject identities in prepared datasets so as to protect privacy and confidentiality is a common condition of both IARPA and IRB approvals of data collection protocols. This would likely extend to data collections involving human subjects administered in an online manner.	No	N/A
26	Data collection by International Participants - Are there any additional restrictions on international partners in collecting part of the researcher-collected datasets?	1.D.1.2, p. 13 6.B.3, p. 49	The U.S. Government has no additional restrictions applicable to international data collection efforts. Offerors are advised to have their international partners verify that they will have no such restrictions imposed by their own governments which may impact a prime Contractor's ability to meet the requirements.	No	N/A
27	Static/UAV Data Collection - Is there any additional guidance on how to split between data collected from static cameras and those from UAVs? i.e., can we have 50% of the volunteers in BRC from static cameras and the other 50% from UAVs or is there a preference for UAV-collected datasets?	1.D.1.2, p. 14	See answer to Question #19.	Yes	1.D.1.2, p. 15
28	It is indicated that the pipeline should migrate to deployment on a UAV for Phase 3. Should we assume the pipeline runs on a single UAV independently versus upon multiple UAVs that are in continual communication. The latter approach would assume some type of comms between UAVs for full fusion of partial fusion results or communication with a ground station. This is an interesting problem but seems outside the scope of what BRIAR is looking for.	1.A.2.3, p. 8 1.F., p. 21	For Phase 3, a version of the integrated BRIAR system will need to operate on edge hardware that can be deployed on an individual field or aerial platform. BRIAR systems running on multiple concurrent platforms in a distributed manner is out of scope.	Yes	1.F, p. 21
29	It appears that there are no restrictions on the unclassified, government provided development data set. Would it be permissible for a researcher at a foreign University to access and download the government provided development dataset to a server that resides at their University?	1.D, p. 10 1.D.3, p. 17	Correct, foreign Performers will be permitted to download, store, and use BRIAR datasets at their sites, subject to an approved Privacy Plan. Government-provided BRS and BTS datasets will be Unclassified and are not expected to be subject to export control. However, they will be provided to Performers under a controlled distribution, limited to BRIAR Program use only.	No	N/A
30	Are you able to provide example images or example sensors that may be used?	1.A.1, p. 4	Example imagery and videos are available in the BRIAR Proposers' Day slides and webinar available on the IARPA website ( <a href="https://www.iarpa.gov/index.php/research-programs/briar/baa">https://www.iarpa.gov/index.php/research-programs/briar/baa</a> ). No additional imagery will be provided prior to program kick-off.  Example sensors are not available at this time. Per Section 1.D.1.2 (p. 15): "For planning purposes, Offerors are to use their best judgement based on the information provided here for their selection of UAV and sensor equipment when planning and proposing their BRCs."	No	N/A
31	Is the body height ~200 pixels and face width of ~20 pixels a good representation of what to expect of a person at 1 km with the sensors that may be used?	1.A.1, p. 4	Yes, those resolution dimensions are approximations of what can be expected in BRIAR imagery captured at 1km range.	No	N/A
32	Are you able to elaborate on the gallery enrollment requirements per phase? What I gather is that phase 1 is a short video or even stills of the face and body in ideal indoor lighting conditions, basically a professional portrait. Phase 2 seems to be outdoor portraits or indoor but the person is not facing the camera or obstructing their face or body. Phase 3 seems to be "small range" so in addition to obstructions and noncompliance they also need to be captured from further away and/or with a small pitch angle.	1.A.2.1, p. 6 1.D.1.2, p. 14	See answers to Questions #16, #19.  Enrollments shall include face imagery consistent with mugshot capture conditions, gait video from multiple views, and still imagery of the whole body. Phase 1 BRC will focus on controlled indoor capture. Phase 2 will add partially controlled enrollment video, and Phase 3 will add small to medium range unconstrained video imagery. Additional guidance and standards will be provided at program kick-off.	Yes	1.D.1.2, p. 14-15
33	The BAA states that collections will occur in a variety of outdoor terrain. There is also mention in 1.A.1 that WB biometric recognition needs to work in an indoor and outdoor setting. Is the expectation that data is collected both indoors and outdoors? If so, does a closed stadium or hangar count as indoors?	1.D.1.1, p. 12	Yes, data will be collected both indoors and outdoor. The majority of indoor data will be collected as part of controlled or partially controlled gallery enrollment data. However, the program may involve unconstrained indoor imagery as probe data depending on transition partner use cases.  Yes, a closed stadium or hangar would count as indoors.	No	N/A
34	What is meant by "subject pose" in the annotations and metadata? Are these lines that mark the limbs or is this more akin to movement classification: walking, standing, etc?	1.D.1.1, p. 13	Subject pose is the positional orientation of the subject's body and/or face in a common frame of reference with respect to the sensor view.	Yes	1.D.1.1, p. 13

35	What specific information will be provided with environmental conditions? Is this location, time of day, sky cover percentage, humidity, temperature, wind speed, indoor/outdoor?	1.D.1.1, p. 13	The current plan is to provide as much detail as possible regarding environmental conditions, to include most if not all of the items listed in the question. More details will be provided at program kick-off.  Per Section 1.D.1 (p. 11): "Proposals must specify the development data needed to carry out the proposed research and what data characteristics are necessary for their approach(es) to be successful at meeting program objectives." BRIAR T&E will work to accommodate reasonable data and annotation needs.	No	N/A
36	What specific information will be provided with atmospheric turbulence conditions? Is this low, medium, high as shown in the BAA figure 2? Or will this be specified by a value such as Cn <sup>2</sup> ?	1.D.1.1, p. 13	For BRS and BTS datasets, both qualitative descriptors (i.e., low, medium, or high) and Cn <sup>2</sup> values will be provided as metadata.	Yes	1.D.1.1, p. 13
37	The BAA mentions enrollments should follow NIST/ANSI standards but there is no document referenced in the BAA. Is this NIST link an acceptable enrollment standard for faces in phase 1? <a href="https://www.nist.gov/system/files/documents/2019/04/23/frvt_quality_concept_1.0.pdf">https://www.nist.gov/system/files/documents/2019/04/23/frvt_quality_concept_1.0.pdf</a>	1.D.1.2, p. 14	Enrollments should be <b>informed</b> by standards and best practices. Relevant standards include ANSI/NIST-ITL 1-2011 (Biometric Data), ISO/IEC DIS 39794-16 (Full Body), ISO/IEC DIS 39794-17 (Gait). At program kick-off, additional guidance will be provided to Performers by the Government for BRIAR enrollment capture to standardize BRC, BRS, and BTS methodologies and image capture conditions.	Yes	1.D.1.1, p. 12 1.D.1.2, p. 14
38	The BRCs require each subject to have two different clothing configurations. Is there a more specific requirement such as changing all visible clothing or is a change of shirt enough? Is one configuration to be used for the gallery enrollment and the other for the probes? Can the clothing configuration of the subject in the enrollment video match the clothing configuration of the subject in a probe video or is this discouraged to prevent simple ReID based on impermanent attributes?	1.D.1.2, p. 14	Clothing changes should be sufficiently different to hinder solely appearance-based re-identification approaches. A minimum of two clothing configurations are required in both probe and gallery imagery, such that a minimum of two cross-clothing biometric pairs can be curated for each subject (e.g., Clothing-A vs. Clothing-B and Clothing-B vs. Clothing-A). Capturing enrollment and field data with the same two clothing sets is acceptable as data curation combinations can be leveraged to emphasize cross-clothing matching.	Yes	1.D.1.1, p. 12 1.D.1.2, p. 14
39	For the BRCs, the threshold of different sensors per subject is 4. Does this threshold apply to the gallery dataset or only the probe dataset? In other words are we required to acquire gallery enrollment ground truth on the same subject using a minimum of 4 different sensors in addition to having to capture the same subject in at least 4 different probe videos produced by 4 different sensors? Can this requirement be met by having the subject appear in 3 probe videos from 3 different sensors and 1 gallery video with yet another different sensor?	1.D.1.1, p. 12 1.D.1.2, p. 14	Sensor variability in the Government and Researcher Collections applies to probe imagery. Gallery enrollments do not require sensor variability, but some research approaches may benefit from incorporating that into their collection plans.	Yes	1.D.1.1, p. 12 1.D.1.2, p. 14
40	What specific information will be provided and needs to be provided with capture sensor or platform characteristics? Is this focal length, aperture, shutter speed or simply whether it is a UAV, elevated camera, or man-portable telephoto?	1.D.1.2, p. 15	Metadata should include, at a minimum, the sensor make, model, and attachments (e.g., lens details); sensor category (i.e., fixed, portable, or UAV); and the UAV platform make and model (if relevant). Offerors should include any other metadata annotations needed to support their proposed research approach.	Yes	1.D.1.2, p. 16
41	What is meant by "Identity and Sighting Instance Labels"? Is the identity the unique subject # and is the sighting instance label simply a way to differentiate between multiple subjects (eg. red box is subject #146)?	1.D.1.2, p. 15	See answer for Question #12	Yes	1.D.1.2, p. 15
42	What data is required to be collected for "Subject Demographic Information"? Is this the subject's age, height, weight, gender, race, and ethnicity? In terms of demographic diversity and body type diversity, is there a specific target metric or minimum?	1.D.1.2, p. 15	At a minimum subject height, weight, age, gender, and race/ethnicity should be captured. For demographic diversity, offerors should use best judgement informed by recent academic publications involving equitable algorithm performance and bias mitigation research. Additional guidance will be provided by IARPA at program kick-off.	Yes	1.D.1.2, p. 16
43	For the T&E data augmentation service, can the data be augmented for both distance and turbulence individually? Or is this the same augmentation just a different choice of words on how we can request data to be augmented?	1.D.1.4, p. 16	Yes, augmentations to transform distance, turbulence, or combinations of both may be requested as part of a Data Augmentation Request. Per Section 1.D.1.4, "Determination of the details of the data augmentation task and the targeted range of simulated conditions will be by collaboration among IARPA, BRIAR T&E, and the Performer."	No	N/A
44	For phase 1 and 2, the probes are expected to be either range driven or pitch driven but not both. Specifically for the pitch driven probes, will the capture range still be provided as a direct linear distance from sensor to subject? Page 23 of the Proposers Day Briefing has the ground distance measured, but I assume the direct linear distance asked for in the BAA would be the diagonal of the triangle and not the ground distance? Is the pitch angle also a constraint on the subject's head pose so that they are not facing less than X degrees at the camera?	1.D.2, p. 17	Yes, the primary measure of range for elevated platforms will be the distance from sensor to subject and not ground distance measurements. Per Section 1.D.2: "Range: The direct linear distance from the optical sensor to subject".  Sensor-view pitch angle is defined as the angle below horizontal of the sensor optics path. The pitch angle of a subject's head certainly factors into the complete face-to-sensor pitch angle, but is a more complicated and changing variable to control or characterize within a given video clip. Sensor-view pitch serves as an overall proxy for the research challenge and is easier to control and measure. Changes in subject head pose will be present in some imagery, both deviations from zero-pitch and/or zero-yaw head position and time-varying head poses within a single sighting. Offeror's proposals should take these factors into account, both in their research approach and in their planning of Researcher Collections.	Yes	1.D.2, p. 18 (footnote)
45	For purposes of the processing speed requirements, can we assume that the reference templates have already been generated from the gallery?	1.F, p. 19-21	Correct, processing speed metrics do not apply to the time required to enroll and create the reference gallery for search tasks. Reference templates and any supporting gallery construction can be assumed to have already been generated. Processing speed applies to end-to-end processing of the probe including, any necessary detection, localization, feature extraction and search.	Yes	1.F, p. 21
46	Will the program metrics (verification, rank retrieval, etc) be evaluated frame-by-frame or can a single result be outputted at the end of the file when all the data is known?	1.F, p. 20	Per Section 1.F, p. 19: "Performance progress toward program metrics will be computed using track-based templates generated from single or multiple video input probes matched against similar reference templates."	No	N/A
47	You mentioned in the Proposer's Day presentation (timestamp 58:10) that a template can be generated using multiple probes. The program metrics seem to hint that only one probe will be used (1:1, 1:N)? Is multiprobe template generation still an area of interest for the program? If so, for real-time in phase 2 and 3, is it expected that there will be multiple synchronized video streams connected to a single computer for processing? And is the processing speed metric the same with multiple streams?	1.A.1, p. 5 1.F, p. 21	Yes, multi-image templates are a program objective. Per Section 1.A.1: "Multi-image templates that leverage video imagery to produce a single biometric encoding for each individual, regardless of the number of imagery files or frames."  For the processing speed metric, per Section 1.F (p. 21): "For the purposes of this BAA, "real time" is defined as a <b>single</b> 1080p video stream..."  The ability to process multiple time-synchronized streaming video feeds is not a requirement of BRIAR.	No	N/A
48	For real-time processing, is the expectation that a template needs to be generated from frame number 1 that then goes through to the matching phase? Is there a certain number of frames that is acceptable to wait for before advancing to the matching phase or are we to determine a go/no-go metric ourselves?	1.F, p. 21	To process streaming video, a delay in returning match results is permissible due to buffering a minimum system-required amount of video content. Since this is dependent on a Performer's specific research approach, Offeror's shall address how their research approach is compatible with streaming video and any assumptions, limitations, or operating conditions.	Yes	1.A.2.2, p. 7-8

49	With the COVID vaccine being distributed, can we assume that the kick-off meeting around August 1st, 2021 and all other PI meetings will be in person?	1.H.1, p. 29	Per Section 1.H.1 (p. 30): "IARPA reserves the right to hold the meeting virtually for logistical or health and safety reasons." A decision on the program kick-off meeting will be made after source selection has concluded. For all program Kick-off and PI Review Meetings, Offerors should plan and budget for in-person meetings in the Washington, DC region.	No	N/A
50	The BAA states there is a 40 page limit not including required attachments. We are able to attach up to three relevant papers and a one-page summary for each but to my understanding these are optional attachments. Do optional attachments count against the 40 page limit?	4.B.1.d, p. 39	Section 4 – Attachments: "The attachments listed below shall be included with the proposal, if applicable, but do not count against the Volume 1-page limit." "Attachment 5: Relevant Papers" is included in the list of Attachments.	No	N/A
51	Can a final solution be based on a system component (e.g., an FR algorithm) that is a licensed, proprietary technology?	1.C, p. 9 4.B.1.d, p. 39	Per Section 1.C, "Approaches that consist merely of integrating currently existing software" are out of scope. However, a proposed research approach can be based on improving and innovating beyond pre-existing research techniques and prototypes.  Regarding IP Rights, per Section 4.B.1.d: "IARPA anticipates that achieving these goals for the BRIAR program may necessitate a minimum of Unlimited Rights in all deliverables. However, there may be any number of other approaches to intellectual property rights to achieve IARPA's program goals."	No	N/A