The BRIAR Dataset: A Comprehensive Whole-Body Biometric Recognition Benchmark at Extreme Distances and Altitudes (Collections 1-6)

Anonymous CVPR submission

Paper ID *****

Abstract

The Biometric Recognition and Identification at Altitude and Range (BRIAR) program aims to extend biometric capabilities into severe operational environments characterized by long ranges, atmospheric turbulence, and elevated viewpoints. This paper introduces the program's two final government collections expanding activities, environments, viewpoints, distances, and modalities, and adding appearance and environmental stressors. We refine curation/evaluation and include sequestered subsets to support controlled assessments. These updates strengthen BRIAR as a comprehensive resource for whole-body, face, and gait recognition at altitude and range, while maintaining privacy-first practices.

1. Introduction

Biometric recognition and identification have achieved notable progress with the advent of deep learning and the steady growth of large-scale datasets. However, performance degrades in operationally challenging scenarios: long standoff distances that reduce facial crops to low resolution; elevated viewpoints from rooftops and unmanned aerial platforms that alter appearance geometry; adverse weather and sensor motion that introduce blur; and natural, unscripted subject behavior that creates occlusions and complex interactions. In these settings, whole-body and gait cues become essential complements to facial identification, yet they remain vulnerable to clothing changes, pose variability, atmosphere/weather condition, and multi-person interactions. Advancing robust identification under these conditions requires datasets that systematically span altitude and range, viewpoints, modalities, activities, and environmental factors, accompanied by rigorous curation, rich metadata, and reproducible evaluation protocols.

The Biometric Recognition and Identification at Altitude and Range (BRIAR) program – funded by Intelligence Advanced Research Projects Activity (IARPA) – addresses these needs and advance the state of the art. The first BRIAR dataset paper[1] introduced two BRIAR Government Collections (BGCs), BGC1 and BGC2, establishing a foundation for long-range and elevated biometric research with controlled and field scenarios, multi-platform ground and aerial sensors, cross-clothing subject capture, and standardized curation and annotation. Building on this foundation, BGC3 and BGC4 added new locations, group activities, a mock urban environment, refined curation/metadata, and formalized evaluation protocols [4].

This paper presents two new collections, BGC5 and BGC6, that substantially broaden the scope, realism, and difficulty of biometric recognition at altitude and range. BGC5 introduces action-oriented field scenarios including toy-weapon aiming, mock underground exchanges, and perimeter checking that elicit diverse poses and occlusions. (See figure 4.) It deploys a larger, more naturalistic mock-city with expanded elevated vantage points and viewing angles, adds two additional cameras at 1,000 m to strengthen extreme-range coverage, and incorporates non-visible spectrum sensors to enable cross-spectral research and fusion. To support controlled assessments, a subset of BGC5 data has been sequestered for additional evaluation. BGC6 complements these expansions by introducing a camouflage military uniform clothing set to reduce distinctive appearance cues and stress whole-body and gait-based identification (Figure 2), and by collecting under high-wind conditions that induce pronounced camera jitter and motion artifacts—conditions to which long-range sensors are particularly susceptible—further challenging detection, tracking, stabilization, and enhancement. Together, these collections extend BRIAR's role as a comprehensive resource for studying identification at altitude and range across modalities, activities, and environments, and for advancing algorithms in cross-spectral fusion, action-aware biometrics, stabilization and jitter mitigation, and appearance-agnostic whole-body and gait recognition. They build on prior

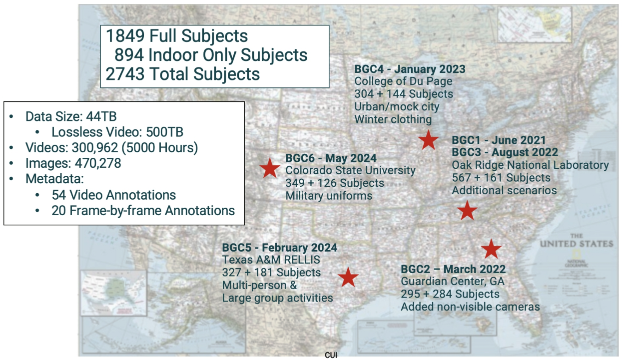

Figure 1. Summary of all BRIAR BGC collections.

BRIAR protocols [1][4] and explore underexamined operational aspects. Figure 1 provides an overview of all BGC collections: locations, subject counts, and data volumes.

Currently, BRIAR datasets are available to U.S. Government (USG) researchers and their contractors upon request. Interested parties should contact the authors, who will forward the request to the appropriate USG representatives.

2. Related Work

The field of biometric recognition has recently pivoted from analyzing high-quality, cooperative imagery to addressing the "operational gap" of identifying non-cooperative subjects in unconstrained environments. Following the initial release of the BRIAR datasets, a couple of public datasets were introduced in 2025 and 2026 to target specific components of this challenge, such as aerial surveillance and low-resolution recognition. However, these benchmarks remain limited in scale, environmental realism, and physical fidelity when compared to the expanded BRIAR corpus.

For instance, recent benchmarks including AG-VPReID[6] and DetReIDX[2] target cross-platform aerial-ground re-identification at altitudes up to 120 meters. While they introduce specific novelties — such as the 3,027 identities within AG-VPReID and the multi-session clothing changes incorporated into DetReIDX — they remain constrained in operational scope. AG-VPReID relies on a highly restricted sensor array of only six cameras, whereas DetReIDX is limited to 509 identities. Furthermore, Their 120 meter maximum ranges preclude severe atmospheric turbulence that degrades biometric signals at tactical ranges. In contrast, the cumulative BRIAR dataset encompasses over 1,800 subjects recorded across multiple years, utilizing a comprehensive array of sensor platforms to capture profound seasonal clothing variations and extreme distances extending to 1,000 meters.

The necessity of such a physics-aware, multi-modal approach was independently validated by the 2025 NIST Face in Video Evaluation (FIVE)[3]. This government evaluation focused on facial recognition performance and utilized sequestered subsets of the BRIAR dataset to test long-range capabilities. The results quantitatively defined a "performance wall" at ranges exceeding 650 meters, where facial features degrade below usable thresholds. This finding reinforces the central hypothesis of the BRIAR program: that at tactical ranges, reliance on facial biometrics alone is insufficient. By providing synchronized Whole-Body (WB) and gait signatures alongside face data, and introducing adversarial elements such as the camouflage protocols and high-wind camera jitter in BGC5 and BGC6, BRIAR provides the only comprehensive benchmark for developing systems capable of maintaining identity persistence when the facial signal collapses.

3. Data Collection Methods

BGC5 and BGC6 retained the broad strokes of the methods and equipment used during the first four collections[1][4]. Given BRIAR's overall goal of bridging the "operational gap" between cooperative and non-cooperative, constrained and unconstrained environments, data collection events consisted of "controlled" (indoor, structured) and "field" (outdoor, less-structured) portions, with participant activities and data collection scenarios running a gamut of operational conditions. Participants were instructed and led through activities by members of the collection team who recorded non-video data, metadata, and activity timestamps using a custom-built proctor application.

The unique characteristics of the collection locations (Figure 1): the terrain, site layout, climate, weather, wildlife, available facilities and infrastructure, local population density and demographics, etc. not only shaped the resulting dataset in terms of imaging conditions and demographics distributions, but also required accommodations and enabled additions to its experimental design that co-determined its composition. The logistical implications of the site-specific factors were wide-ranging and affected every aspect of the data collection. For instance, although field camera and support system placements were primarily determined by a desired distance and view angle, they were co-determined by material handling considerations, power and network availability, participant transportation routes, staging and view-blocking to exclude non-participants from the dataset, among other factors.

3.1. IRB protocol and Recruitment

The Institutional Review Board (IRB) protocol for BRIAR was approved by the Central Department of Energy IRB (CDOEIRB). BRIAR was a multi-site study and required IRB modifications for each collection. The IRB approved consent form included a check box that allowed a participant to decide to be included in public settings (journals, publications, presentations and reports) and only those par-

Make	Model	Notable Specs
Axis	M3057-PLVE	Fixed dome, Panoramic
	P1468-LE	Fixed bullet, 6.2-12.9mm, 4K
	P3245-VE	Fixed dome, 3.9-8.9mm
	P5655-E	PTZ dome, 4.3-137.6mm
	Q1798-LE	Fixed bullet, 12-48mm, 4K
	Q6215-LE	PTZ, 6.7-201mm
Basler	acA2040-90um/uc	Color/mono CMOS
	acA2040-120um/uc	Color/mono CMOS
	acA4112-30um/uc	Color/mono CMOS, 4K
Bosch	MIC-9502-Z30WVF	PTZ, Dual vis & LWIR
Digital Watchdog	DWC-MPTZ336XW	PTZ dome, 4.6-165mm
Fujinon	D60x16.7SR4DE-V23	Motorized zoom lens, 20-1200mm
	SX1600	Fixed bullet, 20-1600mm
Hanwha Techwin	PNP-9200RH	PTZ dome, 4.8-69mm, 4K
	QNP-6250H	PTZ dome, 4.4-111mm
	QNP-6320H	PTZ dome, 4.4-102.2mm
	TNU-4051T	PTZ, 35mm, LWIR
	XNP-6550H	PTZ dome, 4.7-261.4mm
	XNP-9300RW	PTZ dome, 5-150mm, 4K
	XNZ-6320A	Fixed bullet, 4.4-142.8mm
Honeywell	HBW8PR2	Fixed bullet, 3.7-11mm, 4K
	HDZP252DI	PTZ dome, 4.8-120mm
Kowa	LMZ20750AMPDC-XF	Motorized zoom lens, 20-750mm
MOOG/Quickset	MPT-50	PTZ, Dual vis & MWIR
	MPT-90	PTZ, 33.4-2000mm
Nikon	D5600	DSLRs used for controlled images
Sigma	150-600mm DG	DSLR lens, adapted C-mount
	18-300mm	DSLR lens, adapted C-mount
Speco	O2VLB8	Fixed bullet, 2.8-12mm
Swarovski	95mm	Spotting scope, 900-2100mm
	115mm	Spotting scope, 900-2100mm
Tamron	18-400mm	DSLR lens, adapted C-mount
Vivotek	TB9331-E	Fixed bullet, 19mm, LWIR

Table 1. Cameras used during the BGC5-6 data collections. Standalone sensors and lenses were used together in various configurations. Multiple cameras of the same model were often used simultaneously. Approximately 66 camera systems were consistently recording during BGC5, 45 during BGC6.

ticipants that checked the box can have their images and/or videos used in public settings. To protect subjects' privacy, data are stored on properly monitored and protected computer systems. Data use agreements and IRB approval are required for any organization that wants to use the data. Participants filled out a de-identified demographic survey and they received compensation for their participation.

Recruitment success varied by location and various resources were used to reach possible participants. Social media sites, flyers, event calendars, newspapers and word-of-mouth referrals were utilized for recruitment. A website was utilized for sign-ups, which allowed participants to review study details, sign-up for a time slot, reschedule and cancel on their own.

3.2. Controlled Collection

Controlled data collection was purposefully kept as similar as possible to previous BRIAR efforts. Participants completed the entire controlled collection twice, the second time in different clothing; for BGC6, the second set of clothing for each participant was a camouflage-patterned uniform provided by the collection team, as seen in Figure 2. Each individual had a series of still images taken from 0°, 20°, and 30° elevation angles while facing at 45°-increment angles to the cameras: five angles of face images (profile to opposite-profile) and eight angles of whole-body images (a full rotation) were collected. Image sets were captured twice in this way – once with neutral facial expressions and relaxed standing poses, and again with the participant smiling for the face images and standing with arms and legs extended in an "X" pose for the whole-body images. Next, the participant completed two activities in the semicircular area shown in Figure 2 with commercial video surveillance cameras arranged around the perimeter. The participant was instructed to walk back and forth along the diameter line of the semicircle, then to walk randomly within the area for 60 seconds, while talking or texting on a cell phone for the last 30 seconds.

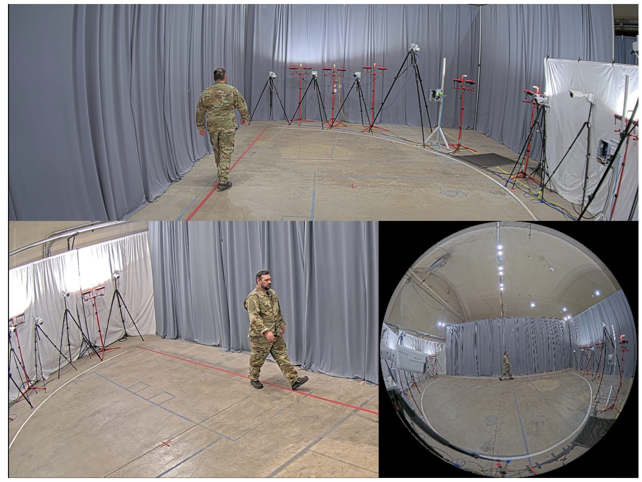

Figure 2. Controlled video frames from BGC6. Note the camouflage-patterned uniform worn by the subject. *Permission granted by subject for use of imagery in publications.*

3.3. Field Collection

As with previous collections, part of the field data collection consisted of activities performed in a ten-by-ten meter square area designated with ground-marking paint; the "field box" seen in the blue cutouts in Figure 3 consisted of a white-bordered square with colored lines positioned at 45°-angles that were used to direct structured activities. A set of close-range cameras and three RGB depth cameras, positioned five to twenty meters from the center of the square, recorded participants at relatively consistent view-angles and distances across collections. At longer ranges, between 100-1000 meters from the field box, the location, make, model, and number of cameras were different for each collection, but were generally clustered in stations with common power, network, and recording infrastructure. Each station marked on the layouts in Figure 3 had at least one operational camera focused on the subject

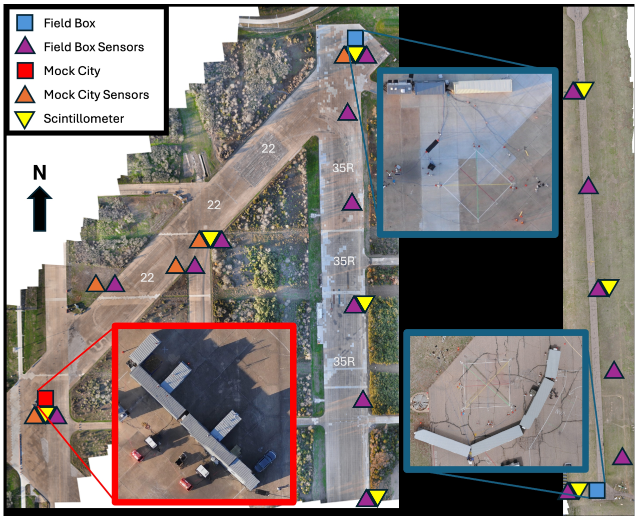

Figure 3. Layout of BGC5 (left) and BGC6 (right). Triangles mark the locations of the camera stations, as well as the scintillometers used to estimate Cn^2, a measure of atmospheric turbulence. The blue-bordered cutouts show the orientation of the field box for each collection, and the red-bordered cutout shows the mock-city layout.

Figure 4. Example frames from scenarios in BGC5's mock-city. Clockwise from the top right: An unstructured group activity with tables displaying snacks and drinks. A subject walking up and down a small set of stairs three times in both directions. A subject burying and digging up objects in a sandbox. A subject pretending to shoot at numbered targets with a toy rifle and pistol. A subject performing a visual inspection of a car, then entering and exiting the driver's side. *Permission granted by subjects for use of imagery in publications.*

area for the majority of the collection. Over the course of the project, the collection team operated a wide variety of camera systems, both custom and commercial-off-the-shelf. Target image quality drove system selection and placement, i.e. HD or better, with optics and/or digital zoom capable of producing an image closely bounding or within the field box. Custom pairings of long-range lenses and industrial CMOS sensors were often utilized at longer ranges. Although not included in the dataset and evaluation protocols established for the performer teams, non-visible spectrum video was recorded from medium- and long-wave infrared sensors during BGC5 and BGC6. A representative list of sensor models used in BGC5 and BGC6 can be seen in Table 1.

The software for the custom cameras was fully rewritten following BGC4. Instead of recording both compressed and raw video to local hard drives, the cameras streamed H.264-compressed video over RTSP to the same network video recorders used by the commercial systems. The custom camera software retained the ability to write raw video to a local hard drive when triggered over the network by the proctor application.

Every participant in the field collection included in the BRIAR dataset completed a "standing" and a "structured walk" activity, in which they were instructed to stand facing along the colored lines of the field box, or walk back and forth along them in a specific order. Each subject also performed additional activities that varied between collections, which are detailed in the following sections.

3.3.1. BGC5

The most complex BRIAR collection, BGC5 was hosted at a former airbase on Texas A&M's RELLIS campus in Bryan, TX during February 2024. BGC5 added a "mock-city" location, constructed out of shipping containers with some modified as portable offices, at which more interactive, realistic scenarios were conducted by participants as described in Figure 4. These activities were recorded by cameras installed throughout the "city" and at range up to a kilometer away. The mock-city was located on the opposite end of the same runway as the normal field collection. Multiple locations and the intersecting runways of the airstrip, seen in Figure 3 presented unique challenges for sensor layout. Camera installations along runway #22 had line-of-sight to both the mock-city and the field box, and as such every installation along the runway was able to support cameras pointing in both directions. Runway #35R provided an alternative view-angle of the field box from several distances; this data was purposefully sequestered as a holdout test set for BRIAR performer team model evaluations.

In addition to the two standard field box activities, each participant completed three more scenarios out of five possible. All involved unstructured walking within the field box: simply walking around, walking while interacting with a cell phone, walking with a backpack, passing a backpack between a group of participants, and passing a large cardboard box between a group of participants. In order to streamline subject flow through the collection, especially given the additional mock-city location, participants did not change clothes during the field collection; as usual, the controlled data features each subject in two sets of clothing.

3.3.2. BGC6

BGC6 was the final BRIAR data collection event. It was also hosted at an inactive airfield – Christman Airfield on Colorado State University, Foothills Campus in Fort Collins, CO. Data was collected from the end of April through May 2024. The ground was slightly bowed from one end of the runway to the other, significantly limiting visibility at range. Scaffolding was built at the stations 750 and 1000 meters from the field box to elevate the cameras high enough to view the participants over the hump in the center of the runway. Extendable masts were also used at these two stations, each with a single camera and tensioned guylines to keep the cameras stable. Frequent high winds at the site caused significant camera shake, especially for the mast- and scaffold-mounted systems.

Participant image and video data were collected at both the field and controlled sites in two sets of clothing: in the clothes the subjects arrived in, and in camouflage uniforms provided by the collection team. Similarly to BGC5, subjects participated in the two consistent standing and structured-walking activities, and also performed two activities selected at the beginning of each day out of four possible: the individual cell-phone, individual backpack, group backpack, and group box-pass scenarios mentioned in the section above.

3.4. UAV Collection

BRIAR's aerial operations were extensive, with up to four aircraft flying concurrently. Safety and airspace coordination were central: site selection balanced the ground-range requirement with sufficient distance from active airports (e.g., BGC2 near Perry–Houston County Airport). To mitigate the risks of operating in close proximity to active airports, Notices to Airmen (NOTAMs) were issued, routine coordination with local Air Traffic Control (ATC) was maintained, aviation-band radios were available to contact manned aircraft, multiple visual observers were employed, and contingency plans were enacted for rapid diversion or recovery. Operations were conducted under either Part 107 or Certificates of Authorization (COAs).

To overcome outdated commercial imagery, an eBee mapping Unmanned Aerial System (UAS) was flown on scouting missions to produce approximately 1 cm Ground Sampling Distance (GSD) orthomosaics. Large-format "battle boards" enabled precise placement of ground sensors, flight-path planning, and situational awareness. Near the end of deployment, the eBee was re-flown to geolocate installed ground sensors and compute exact ranges. Figure 3 exemplifies these mapping efforts.

Table 2 shows a summary of UAS platforms used for the collections. Weather — especially wind — was a persistent constraint. Small UAS are sensitive to precipitation, icing, and high winds; for example, BGC6 in Fort Collins frequently grounded multirotors. Where feasible, a tethered FLIR R80D Skyraider (Elistair Safe T2) extended endurance in higher winds. Icing risks (e.g., pitot blockage, control-surface impairment) and temperature/humidity extremes impacted battery performance and crew endurance, prompting conservative sortie planning.

RF/GPS interference and hardware wear introduced communication issues. The R80D occasionally experienced tether power adapter failures and antenna contact wear within the landing legs, leading to intermittent link recognition and preflight faults. At times, the Autel EVO II link appeared to interfere with the R80D, mitigated by temporal/spatial separation.

Despite these challenges, the aerial effort delivered synchronized elevated-view video that complemented ground sensors across altitudes and geometries. High-resolution mapping improved planning fidelity and provided authoritative geolocation, while the mix of fixed-wing, hybrid VTOL, tethered, and small multirotors ensured coverage across varied conditions.

4. Dataset Summary

The BGC5 and BGC6 combined dataset consists of over 300,000 images and over 120,000 videos. Videos range from around 30 seconds up to over 3 minutes in length. See section 3 for image and video collection details.

Figure 5 shows the distributions of subject-to-sensor pitch and yaw angles across controlled images. Figure 6 shows the distributions of distance, elevation, and measured Cn^2 values for outdoor videos collected from non-aerial sensors. Here, C_n^2 is the refractive index structure constant, a measure of the strength of optical turbulence in the atmosphere. Table 3 shows the distribution of modality (face vs whole body) across all videos. These modality labels indicate whether the camera is configured to capture only a subject's face or to capture their whole body in a typical collection event from the given operating distance. Figure 7 displays distributions of various per-subject demographics.

5. Data Curation Methodology

5.1. Curation Process Overview

The data curation workflow has been optimized during each of the collection curation efforts; however, the core steps remain consistent. During BRIAR collections, proctors record timestamps for each subject performing a specific activity in a unique clothing set. Timestamps are utilized to segment individual videos out of full day recordings as well as correlate weather and atmospheric measurements for each moment in time. Subjects are photographed in two unique clothing sets across a variety of poses to serve as controlled imagery for enrollments. Once data is returned

Platform (Model)	Manufacturer	Type / Notable Specs	Camera Payload(s)
Supervolo	Hybrid Project	Hybrid VTOL fixed-wing; 9 ft wingspan; 2-stroke gasoline engine with mid-air restart; ˜100 mph cruise; up to ˜7 hr endurance; Silvus Streamcaster 4400/4200 data links; STANAG 4609 metadata	Trillium HD25 (early) / HD40 (later) gimbals
R80D Skyraider	Teledyne FLIR (Aeryon)	Electric quadcopter; ˜25 min per battery; tether-capable (Elistair Safe T2) for extended endurance; 900 MHz C2 link	SR-EO/IR Mk II (thermal, no zoom); HDZoom30 (EO)
Anafi USA	Parrot	Small quadcopter; ˜30 min battery; DIU Blue / NDAA compliant; integrated payload	Integrated EO + thermal camera
EVO II	Autel Robotics	Small quadcopter; ˜30 min battery; non-NDAA/Blue; integrated payload	Integrated 1” CMOS EO + thermal
RQ-20 Puma	AeroVironment	Fixed-wing, hand/catapult launch; deep-stall landing; 2–3 hr endurance	Mantis i45 EO/IR/thermal gimbal
Edge 130	FlightWave	Electric VTOL tricopter; ˜20 min battery	NextVision Nighthawk2 gimbal
eBee (mapping UAS)	senseFly /AgEagle	Fixed-wing mapping platform; produced ˜1 cm GSD orthomosaics for planning and geolocation	Mapping RGB camera (orthomosaic generation)

Table 2. Summary of UAS platforms and sensor payloads used in BRIAR aerial collections.

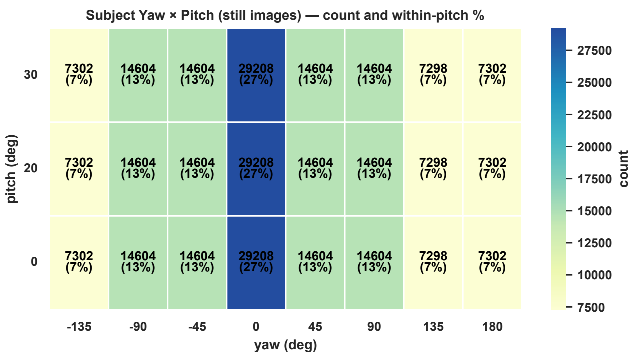

Figure 5. Yaw and pitch angles of sensor relative to subject in controlled images.

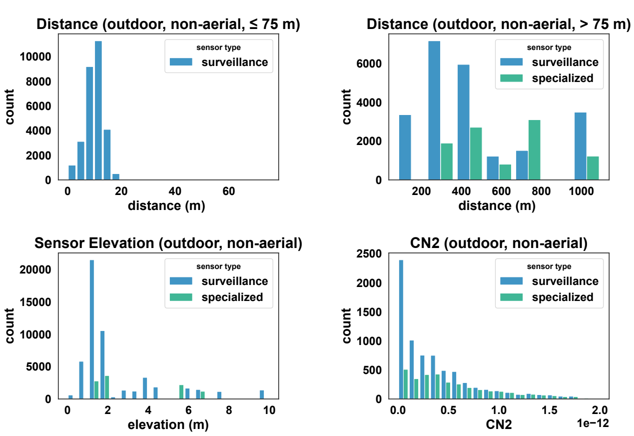

Figure 6. Distance, elevation, and Cn^2 distributions of outdoor, non-aerial videos.

Range	Face	Whole Body
0-10m	4121 (39%)	70731 (65%)
10-20m	1648 (15%)	10752 (10%)
20-200m	0 (0%)	3360 (3%)
200-400m	1897 (18%)	7167 (7%)
400-800m	2588 (24%)	12722 (12%)
800-1100m	435 (4%)	4279 (4%)

Table 3. Video counts and within-modality percentages across distance ranges.

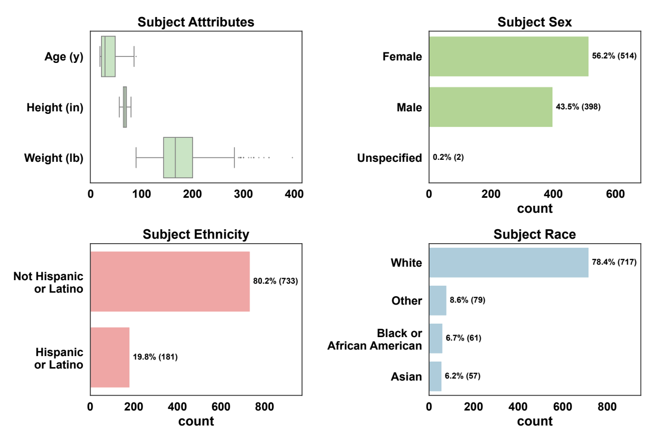

Figure 7. Subject demographics distributions.

from a collection event it undergoes a series of quality assurance tests to check for mislabeled subjects, camera misalignments, timestamp inconsistencies, etc.

The metadata structure remains consistent in the newest datasets such that each unique scenario (subject, clothing set, activity, sensor) is associated with an XML file containing sensor specifications, atmospheric measurements, weather conditions, and subject demographics. After the base XML files are generated and validated they are used as the input to the automated annotation process. The automatic annotations process takes input videos and base XML files and designates person detections, body bounding box coordinates, face bounding box coordinates, and track-level subject designations. This process uses pre-trained open-source models. Manual annotation is done after the automatic annotation process to validate the accuracy of annotations and correct inconsistencies. Finally, the base XML files, automatic annotations, and manual annotations are combined into the final data files which accompany the video files in the released dataset. See [1] and [4] for more detailed descriptions of the curation steps only described at a high level in this section.

5.2. Unified Metadata Curation and QA Pipeline

Previously, each sensor category (surveillance, specialized, and UAV) required a distinct processing pipeline due to heterogeneous recording formats. However, as mentioned in section 3.3, specialized sensors used in BGC5 and BGC6 were standardized to RTSP streaming which enabled consolidation into a single metadata curation pipeline. A centralized metadata server now stores raw data, cleaned data, raw / intermediate / finalized media paths, and detailed operation logs which include configuration parameters. This standardization facilitates additional automated QA checks, more efficient manual reviews, and increased robustness throughout the pipeline.

Centralizing BGC metadata enabled development of more in-depth QA tools which in turn facilitated a faster and more robust curation process. As described in [4] the QA was done to identify time drift, mislabeled subjects, incorrectly attributed metadata, and timestamp collisions. A synchronized multi-view player was developed to view each scenario (unique subject, clothing set, and activity) from all sensors at once for quick identification of misaligned sensors, mislabeled subjects, time drift errors, etc. The other notable tool supported easier time drift correction by allowing the user to view a raw video with a misaligned timestamp side-by-side with a selected subject activity clip. This interface allowed for alignment of problematic scenario clips against a known good sensor and automatic calculation of time-drift correction offsets. Additional automated checks were added to the curation pipeline after the unification was completed but these tools made it possible to more thoroughly address issues which required manual intervention.

5.3. Quality-based Curation Triage

The massive volume of the BRIAR dataset — exceeding tens of thousands of videos per collection — precludes exhaustive manual review for quality check. Operational challenges, such as strong winds inducing severe jitter or exposure failures, necessitate an automated triage system to prioritize human intervention. We developed a suite of metrics to generate a prioritized list of potentially defective videos, allowing the curation team to focus on low-scoring outliers. This hybrid workflow ensures high dataset integrity while providing critical metadata for subsequent covariate analysis of biometric algorithm performance. All metrics are summarized at the video level (mean, std, min/max) to identify both persistent and intermittent quality issues.

To address blur - either by focus or by motion - issues, we developed Peak Mean High-pass Power Spectral Density (PMHPSD). By partitioning frames into grids and isolating the peak high-frequency energy [5], we identify localized sharpness even in complex scenes, effectively flagging videos out of focus or shot at irrelevant areas like the sky. As illustrated in Figure 8a and 8b, sharp frames yield high PMHPSD scores, whereas frames blurred by defocus or camera misdirection result in significantly lower values.

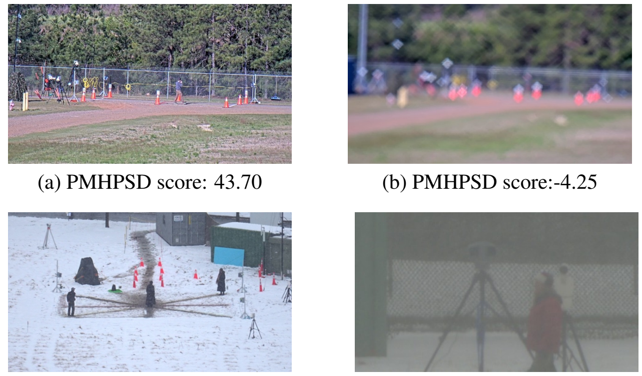

(a) PMHPSD score: 43.70 (b) PMHPSD score:-4.25

(c) AIDI score: 2.66 (d) AIDI score: 0.76

Figure 8. Comparison of Video Quality Metrics for Curation Triage. *Permission granted by subjects for use of imagery in publications.*

For contrast issues, we utilize a novel Approximated Intensity Derivative Integral (AIDI). Unlike standard Shannon Entropy [7] — which is a global measure vulnerable to sensor noise, AIDI sums intensity "slopes" between specific percentiles (several points from 5th to 95th) to better capture structured transitions necessary for identification. As shown in Figure 8c and 8d, a high AIDI score indicates the structured transitions necessary for feature extraction, while a low score reflects the flat profile characteristic of poor exposure.

Finally, we utilize the VidStab framework[8] to quantify video jitter caused by platform instability or extreme weather, flagging sequences where severe motion might hinder subject detection and tracking.

6. Evaluation Protocol Design

BRIAR evaluations use data from BGC1–BGC6 with two probe partitions: FaceIncluded, where automated pose/face estimation indicates a visible head height of at least 20 pixels, and FaceRestricted, where faces are either occluded or with less than 20 pixels in head height, or are not present. Probes are drawn from outdoor semi-controlled and unconstrained recordings spanning close-range ground and elevated sensors, long-range ground sensors (10–1,000 m), and UAV viewpoints, with emphasis on extended ranges and elevated pitch angles (0°–50°). BGC6 introduces a cross-attire condition in which probes use camouflage while galleries remain in civilian clothing. The evaluation gallery sets contain data for distractor subjects who are not part of the probes to simulate a larger gallery.

Evaluations focused on the data collected at longer ranges or elevated views with extreme pitch values. The field data from sensors at 75 m -500 m range and at mast around the field collection box were also incorporated into the gallery. BRIAR evaluations utilize two types of galleries, Simple and Blended, to capture the essence of realistic scenarios. Simple gallery comprises all face and WB images and half of the structured walk sequences from all cameras in the controlled collection. The Blended gallery is derived from the Simple gallery to reflect operational variability: 40% of identities retain Simple-only media, while the remaining 60% are distributed equally across three compositions:

- Mugshot plus close-range surveillance from controlled and field collections (face crops across yaw angles; structured-walk clips excluding toward-camera views).
- Entrance-like close-range surveillance using toward-camera walk sequences.
- Elevated/long-range media drawn from close-range elevated and ground field feeds at 75 m –500 m.

7. Conclusions and Future Work

This paper presented BGC5 and BGC6, extending BRIAR toward more operationally realistic and challenging settings. The new collections introduce new action-oriented mock-city scenarios, more sensors in long standoff distances, elevated viewpoints with camera motion due to high winds, and a camouflage clothing set intended to reduce distinctive appearance cues. In parallel, a unified RTSP-based curation and quality-assurance pipeline with automated triage for blur, contrast, and jitter, together with sequestered subsets and a refined evaluation protocol, support more systematic studies of face, whole-body, and gait recognition at altitude and range.

Several directions appear promising for future work. Although non-visible spectrum video was recorded in BGC5–BGC6, it has not yet been incorporated into public evaluation protocols; integrating these modalities and conducting dedicated nighttime collections under appropriate safety, staffing, and site controls would enable characterization of cross-spectral fusion and low-light performance when visible-spectrum evidence is limited. Seasonal and clothing effects remain an open challenge, as suggested by the BGC4 winter data; systematic studies across layered clothing, bulk, and footwear, potentially including thermal signatures, could clarify degradation mechanisms and inform model design and training. On the algorithmic side, joint modeling of face, whole-body, gait, and action cues, together with stabilization and jitter-mitigation methods and quality-aware fusion, could be evaluated using the new conditions in BGC5–BGC6, with the curation triage metrics serving as covariates for principled analysis.

Pursuing these directions is expected to further enhance BRIAR's usefulness as a benchmark for multi-cue identification in settings where facial evidence may be limited and to guide future sensing, data collection, and evaluation.

8. Acknowledgments

This research is based upon work supported by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via D20202007300010. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

This research used resources from the Knowledge Discovery Infrastructure housed at Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

Notice: This manuscript has been authored by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the US Department of Energy (DOE). The U.S. government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for U.S. government purposes. DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (http://energy.gov/downloads/doe-public-access-plan).

The authors of this paper also wish to acknowledge the contributions of: Knowledge Discovery Infrastructure (KDI) Staff: Dallas Sacca and Ryan Tipton. Proctors and drivers: Raymond Borges-Hink, Nancy Engle, Dale Hensley, Nikki Jones, Michael Jones, Marilyn Langston, Matt Love, Amanda Mottern, Linda Paschal, Christina Peshoff, Donna Pierce, Ryan Styles, and Lauren Torkelson. UAS Pilots: Joe Baldwin, Kase Clapp, Dakota Haldeman, Andrew Harter, Amanda Killingsworth, Matt Larson, Genevieve Martin, Aaron O'Toole, Jason Richards, Brad Stinson, and Brandon Stockwell. Health and Safety Support Staff: Margaret Smith and Miranda Liner.

References

[1] David Cornett, Joel Brogan, Nell Barber, Deniz Aykac, Seth Baird, Nick Burchfield, Carl Dukes, Andrew Duncan, Regina Ferrell, Jim Goddard, Gavin Jager, Matt Larson, Bart Murphy, Christi Johnson, Ian Shelley, Nisha Srinivas, Brandon Stockwell, Leanne Thompson, Matt Yohe, Robert Zhang, Scott Dolvin, Hector J. Santos-Villalobos, and David S. Bolme. Expanding accurate person recognition to new altitudes and ranges: The BRIAR dataset. In *2023 IEEE/CVF Winter Conference on Applications of Computer Vision Workshops (WACVW)*, pages 593–602, 2023. 1, 2, 6

[2] Kailash Hambarde, Nzakiese Mbongo, Pavan Kumar MP, Satish Mekewad, Carolina Fernandes, Gökhan Silahtaroğlu, Alice Nithya, Pawan Wasnik, MD. Rashidunnabi, Pranita

Samale, and Hugo Proença. DetReIDX: A stress-test dataset for real-world UAV-based person recognition. In *1st Workshop on VReID-XFD, Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) Workshops*, 2026. 2

[3] Austin Hom, Patric Grother, and Mei Ngan. NIST face in video evaluation (FIVE), 2025. International Face Performance Conference (IFPC) Presentation. 2

[4] Gavin Jager, David Cornett III, Gavin Glenn, Deniz Aykac, Christi Johnson, Robert Zhang, Ryan Shivers, David Bolme, Laura Davies, Scott Dolvin, Nell Barber, Joel Brogan, Nick Burchfield, Carl Dukes, Andrew Duncan, Regina Ferrell, Austin Garrett, Jim Goddard, Jairus Hines, Bart Murphy, Sean Pharris, Brandon Stockwell, Leanne Thompson, and Matthew Yohe. Expanding on the BRIAR dataset: A comprehensive whole body biometric recognition resource at extreme distances and real-world scenarios (collections 1-4). In *2025 19th International Conference on Automatic Face and Gesture Recognition (FG)*, 2025. 1, 2, 6, 7

[5] Renting Liu, Zhaorong Li, and Jiaya Jia. Image partial blur detection and classification. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2008. 7

[6] Huy Nguyen, Kien Nguyen, Akila Pemasiri, Feng Liu, Sridha Sridharan, and Clinton Fookes. AG-VPReID: A challenging large-scale benchmark for aerial-ground video-based person re-identification. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. 2

[7] C.E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423, 1948. 7

[8] A. Spannbauer. VidStab: Python video stabilization using OpenCV, 2018. 7