Syed Ibrahim Mustafa Shah Bukhari

Blacksburg, VA | +1 (570) 982-2391 | simsb@vt.edu | LinkedIn | Google Scholar | simsb-99.github.io

Education

Virginia Tech **August 2023 – May 2028**

Ph.D. in Computer Science

Research areas: Human-Computer Interaction, Usable Privacy and Security, AI Alignment and Evaluation, and Extended Reality (XR)

Virginia Tech **August 2023 – May 2026**

M.S. in Computer Science · *GPA: 4.00*

LUMS University **August 2018 – June 2022**

B.S. in Computer Science · Minor in Psychology · *Major GPA: 3.92*

Research Experience

Graduate Research Assistant **August 2023 – Present**

Virginia Tech Blacksburg, VA

Advisors: Dr. Brendan David-John and Dr. Bo Ji

PriVisE-Bench: Evaluating Human-Grounded Contextual Privacy Reasoning of Vision-Language Models (VLMs)

- Identified that existing visual-privacy benchmarks model privacy as sensitive-object detection, failing to measure the value alignment of VLMs' privacy judgments with human contextual privacy expectations.
- Reframed scene privacy as a human-grounded contextual judgment and operationalized the evaluation target through privacy label prediction, reasoning alignment, and robustness to prompt wording and ordering.
- Built PriVisE-Bench by mapping a human-grounded privacy taxonomy to 5,800 candidate images and retained 1,069 expert validated whole-scene images across 8 privacy rationales.
- Executed a large-scale evaluation of 14 open- and closed-source VLMs across 326,788 zero-shot, few-shot, and robustness trials using Holm-corrected statistical tests.

Impact: First-author manuscript under review at AAAI'27 Special Track on AI Alignment.

EvaluatAR: Reproducible Cross-Device Evaluation of Bystander Privacy-Enhancing Technologies (PETs) for AR

- Pinpointed that conventional bystander PETs' evaluations rely on live participants, incur substantial setup overhead, and remain device-specific, undermining reproducibility across headsets and slowing iterative development.
- Established a controlled early-stage evaluation protocol that enables researchers and practitioners to reliably assess system feasibility, robustness, and privacy-performance trade-offs under repeatable conditions.
- Designed and implemented EvaluatAR, a modular Unity- and Python-based framework for reproducible, device-agnostic integration, debugging, and testing of real-time AR systems.
- Built standardized sensor I/O hooks, headset-pose alignment, elapsed-time record-replay, runtime logging, and Python-based visualization to abstract device-specific APIs across HoloLens 2, Magic Leap 2, and Meta Quest 3.
- Validated EvaluatAR across 2 PETs and 3 headsets, enabling one experimenter to execute 150 controlled trials in about 60 minutes and rapidly diagnose failures and validate system modifications.

Impact: First-author paper at PoPETs'26; artifact Available and Functional badges.

Rethinking Privacy Indicators in Extended Reality (XR)

- Recognized that conventional XR privacy indicators assume bystanders are attentive and able to perceive visual cues, leaving distracted or otherwise situationally impaired bystanders unaware of nearby sensing and recording.
- Formulated an inclusive design objective to make XR privacy indicators perceivable to situationally impaired bystanders by adapting their modality, salience, and activation to the surrounding context.
- Led 4 participatory focus groups ($N = 8$), eliciting 7 concepts and consolidating them into 5 novel privacy-indicator designs spanning visual, auditory, haptic, proximity-, gaze-, and activity-triggered cues.
- Assessed the proposed designs through an 8-scenario user study ($N = 7$), finding all to exceed the accepted SUS threshold and multimodal indicators being preferred over visual-only cues in privacy-sensitive impairment scenarios.

Impact: First-author paper at ISMAR-Adjunct'25; CCI Research Paper Showcase'26.

"Just Stop Doing Everything for Now!": Security Attacks in Collaborative MR

- Identified unstudied security risks in remote collaborative MR and implemented click-redirection and spatial-occlusion attacks in a networked Unity prototype running on Microsoft HoloLens 2.
- Devised and refined a counterbalanced mixed-methods protocol through 2 pilot rounds, calibrating the collaborative task, attack conditions, behavioral measures, questionnaires, and interview guide.
- Orchestrated the IRB-approved study with 20 participants in 10 remote pairs, capturing task performance, workload, usability, security perceptions, behavioral recordings, and interviews across baseline and attack conditions.

- Triangulated quantitative, observational, and qualitative evidence, finding click redirection increased completion time from 5.25 to 8.54 minutes ($p = 0.02$), while 67% of spatial-occlusion participants misreported their mitigation.

Impact: Published at IEEE VR'25; featured by Computerworld.

Visceral Interfaces for Privacy Awareness of Eye Tracking in VR

- Designed a counterbalanced mixed-methods experiment evaluating 3 visual eye-tracking interfaces on Meta Quest Pro using eye-tracking calibration, visual stimuli, behavioral measures, surveys, SSQ screening, and interviews.
- Conducted IRB-approved study with 40 participants, enabling comparison of privacy attitudes, usability, and interface preferences.
- Synthesized quantitative and qualitative findings into design recommendations, finding that visceral interfaces significantly increased caution toward sharing gaze data with employers and VR entertainment applications ($p < 0.001$).

Impact: Published at ISMAR'24.

Undergraduate Research Assistant **January 2022 – May 2023**

LUMS University Lahore, Pakistan

Advisors: Dr. Muhammad Fareed Zaffar and Dr. Murtaza Taj

TaxData: Human-Centered Remote-Sensing Support for Property Tax Surveying

- Identified slow, manual property surveying as a barrier to tax assessment; curated training, validation, and test data from 30,000 Google Earth images, including 3,000 manually labeled constructed and non-constructed examples.
- Trained deep-learning classification and plot-segmentation models on the curated remote-sensing dataset, achieving 92% accuracy for constructed-versus-non-constructed image classification.
- Translated the model into a full-stack decision-support application that enabled tax officials to define regions of interest and inspect current predictions, reducing property-survey time by 70%.

Impact: Featured in International Growth Centre (IGC March 2022) technical report.

Publications

- "PriVisE-Bench: Benchmarking Human-Grounded Contextual Privacy Reasoning of Vision Language Models over Whole-Scene Images." **S. Bukhari**, A. Chinnala, B. Ji, and B. David-John. *AAAI 2027 Special Track on AI Alignment* (under review).
- "EvaluatAR: A Cross-Device Evaluation Framework for Rapid Prototyping of Bystander PETs in AR." **S. Bukhari**, M. Corbett, B. Ji, and B. David-John. *Proceedings on Privacy Enhancing Technologies (PoPETs 2026(4)).*
- "Rethinking Privacy Indicators in Extended Reality: Multimodal Design for Situationally Impaired Bystanders." **S. Bukhari**, M. Sajid, B. Ji, and B. David-John. *2025 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct 2025).*
- "Just Stop Doing Everything for Now!": Understanding Security Attacks in Remote Collaborative Mixed Reality." M. Sajid, **S. Bukhari**, B. Ji, and B. David-John. *2025 IEEE Conference on Virtual Reality and 3D User Interfaces (IEEE VR 2025).*
- "Visceral Interfaces for Privacy Awareness of Eye Tracking in VR." G. Ramirez-Saffy, P. Chelladurai, A. Vargas, **S. Bukhari**, E. Selinger, S. Foster, B. Heller, and B. David-John. *2024 IEEE International Symposium on Mixed and Augmented Reality (ISMAR 2024).*

Industry Experience

Full-Stack Software Engineer **March 2023 – August 2023**

Senarios

- Architected a cross-platform real-time contact-sharing application, authored design documents, and integrated decentralized components; launched on Google Play Store with 1,000+ downloads.
- Advised 3 clients on Web3-based data immutability while coordinating a cross-functional Agile team.

Backend Software Engineer **August 2022 – March 2023**

GoSaaS Labs

- Developed 7 cloud-to-network-drive features, improving data retrieval speed by 15%.
- Optimized SQL queries to reduce backend response time by 25%.

Teaching Experience

Graduate Teaching Assistant **January 2025 – Present**

Virginia Tech Blacksburg, VA

- Courses: System and Software Security, Professionalism in Computing, and Introduction to Programming.

Undergraduate Teaching Assistant **September 2020 – August 2022**
LUMS University

- Courses: Intro to Programming, AI, SWE, Computer Vision Fundamentals, and Computational Problem Solving.

Leadership Roles

- **Vice President, Virginia Tech Computer Science Graduate Council** **August 2026 – Present**
  Supervise ongoing operations across 4 council departments, coordinating officers and cross-department priorities to support graduate-student programming and council initiatives.
- **Treasurer, Virginia Tech Computer Science Graduate Council** **August 2025 – May 2026**
  Identified and secured an alternate funding source, enabling the council to fund 44 graduate-student events with over $5,000 in support.
- **Webmaster, Virginia Tech Computer Science Graduate Council** **August 2024 – May 2025**
  Managed and maintained the council website, keeping graduate-student resources, event information, and council updates current and accessible.
- **iXR Ideathon Organizing Committee** **February 2026 – April 2026**
  Managed logistics for a 2-day event planned for 40+ participants and coordinated allocation of a $2,500 event budget.

Service & Awards

- **Mentoring:** Mentored Kim Giordano and Amulya Chinnala; UR2PhD-certified Graduate Student Mentor.
- **Reviewing:** *VRST'26*, and GenXR-ID workshop at *CVPR'26*.
- **Awards:** Tapia Scholarship 2026, CCI Cyber Innovation Scholar (2 awards; $3,500 total), PoPETs'26 stipend ($840), and Dean's Honor List.

Technical Skills

Programming Languages: Python, R, C++, C#, SQL, JavaScript
Tools and Libraries: Unity, OpenXR, MRTK, Android XR, Meta XR SDK, Meta Quest 3, HoloLens 2, Magic Leap 2, Galaxy XR, PyTorch, TensorFlow, Scikit-learn, OpenCV, Pandas, NumPy, Matplotlib, Ollama, Node.js, Express.js, Git, Docker, Microsoft Azure, Amazon Web Services, Linux, SSH, HPC clusters
Research Methods: Mixed methods, controlled experiments, within-subjects designs, counterbalancing, participatory design, focus groups, usability testing, qualitative and quantitative analysis, statistical analysis, human-grounded evaluation, multimodal reasoning, benchmark design, prompt engineering, zero-/few-shot evaluation, prompt robustness, AI alignment evaluation, deep learning, AR/VR/MR prototyping