SAIFULLAH

PHD CANDIDATE AT DEUTSCHES FORSCHUNGSZENTRUM FÜR KÜNSTLICHE INTELLIGENZ (DFKI) GMBH

☐ (+49) 0162-8893391 | ✉ saifullah3396@gmail.com | Personal Website | 🎓Google Scholar ID | ⌨ saifullah3396 | 🔗 saifullah3396

Software engineer with 7+ years of experience in research and development of end-to-end software pipelines for AI, deep learning, computer vision and robotics.

Education

RPTU University of Kaiserslautern-Landau
Kaiserslautern, Germany

DOCTORATE (PHD) IN COMPUTER SCIENCE (CS)
July 2021 — ongoing

- Thesis Title: DocXSR: Towards Explainable, Secure, and Robust Document AI
- Designed and developed DocXSR, a framework to strengthen the trustworthiness and regulatory compliance of existing Document AI systems through three core pillars: explainability, security, and robustness.

National University of Sciences and Technology (NUST)
Islamabad, Pakistan

MASTER OF SCIENCE (MS) IN ROBOTICS & INTELLIGENT MACHINE ENGINEERING (RIME) · GPA: 1.0
September 2016 — June 2019

- Thesis Title: Development of a Generalized Whole-Body Motion Architecture for a Humanoid Robot
- Developed a generalized whole-body motion architecture for the NAO humanoid robot that enabled dynamic motion planning of various tasks. Its capabilities were showcased through the design of an improved omni-directional walking engine, an optimized kicking engine, and an adaptable targeted diving motion deployed on real robots.
- Thesis link: `https://tinyurl.com/ms-thesis`

National University of Sciences and Technology (NUST)
Islamabad, Pakistan

BACHELORS (BS) IN MECHANICAL ENGINEERING · GPA: 1.31
September 2012 — June 2016

- Thesis Title: Dynamically Stable and Impact-Controlled Humanoid Kick Engine
- Developed an omni-directional kick engine for the NAO humanoid robot based on momentum control for accurate ball delivery to a desired location.

Experience

German Research Centre for Artificial Intelligence (DFKI)
Kaiserslautern, Germany

PHD RESEARCHER
August 2021 — ongoing

- Developed domain-tailored explainable deep learning methods, including interpretable neural networks, model-agnostic XAI approaches, and visual counterfactual explanation methods, to enhance transparency and trust in both uni-modal and multi-modal document AI. (See publications [3], [6], [8], [9], [12])
- Developed methods to mitigate privacy risks in existing deep learning–based document AI models using differential privacy, federated learning, and private synthetic data generation. (See publications [5] [6][7][13])
- Developed methods to improve the robustness and scalability of existing document AI systems by identifying model failure modes against unseen data and reducing data redundancy. (See publications [1] [2][10][11])
- Developed tools for efficient evaluation and analysis of the XAI methods in multi-modal document understanding (see `https://github.com/saifullah3396/torchxai`).
- Designed training and inference pipelines for document AI, with integrated monitoring and experiment tracking, based on modern frameworks including PyTorch, PyTorch Ignite, MLflow, Ray, Hydra, and LakeFS (see `https://github.com/saifullah3396/atria`).
- Deployed large-scale research experimentation pipelines on the SLURM cluster in distributed setups.
- Engineered and deployed end-to-end document AI solutions, including microservice-based backends and web frontends, using FastAPI, PostgreSQL, Docker, and modern JavaScript frameworks (React, Next.js).

Intelligent Robotics Lab (IRL), National Center of Artificial Intelligence (NCAI)
Islamabad, Pakistan

TECHNICAL LEAD
February 2021 — July 2021

- Developed end-to-end software pipelines for two projects: (1) detection and tracking of people, and (2) detection and tracking of faces and text on live news channels. These projects involved optimizing deep learning models using TensorRT (C/C++) and deploying them on edge devices with NVIDIA DeepStream SDK. Additionally, developed middleware based on Apache Spark and MongoDB for real-time analytics.

RESEARCH ASSOCIATE   *July 2019 — February 2021*
- Developed a web server backend using Python/Django, along with an admin dashboard frontend built in React JS for a realtime pedestrian tracking.
- Worked on deploying a semi-humanoid Pepper robot in a banking environment to provide interactive support sessions to users (Android).
- Led the development of a ROS-based SLAM pipeline for UGVs to enable complete autonomous behavior of Kobuki robots.
- Led the development of an integrated software pipeline (ROSPY, TensorFlow, and OpenAI Gym) for deploying reinforcement learning algorithms on unmanned aerial vehicles (UAVs) in dynamic environments.

Ingenio   *Islamabad, Pakistan*
TEAM LEAD ROBOTICS (PART-TIME)   *March 2020 — September 2020*
- Led the development of the Fortnite Automation product, aimed at creating a robot-controlled player in Fortnite PS4 using an external controller and live AI-based object tracking on edge devices.
    - Developed a USB interface device for overriding PS4 controls through an external Jetson Nano.
    - Deployed Tiny-YoloV3 with NVIDIA DeepStream SDK on NVIDIA Jetson Nano to detect and track targets.
    - Deployed a ROS-based framework for autonomous control of player behavior, such as movement, target tracking, and shooting.

Robotics & Intelligent Systems Engineering (RISE) Lab, NUST   *Islamabad, Pakistan*
TEAM LEAD - 2018-19, TECH LEAD 2017, PROJECT ROBOCUP (PART-TIME)   *September 2017 — June 2019*
- Team-NUST affiliated with RISE Lab, NUST, was established in 2013 with the aim of carrying out research in the domain of humanoid robotics and to participate in the Standard Platform League (SPL) of the annual international robotics competition, RoboCup.
- Design and development:
    - Base software architecture for the Nao robot used in all further research on the robot.
    - Vision modules for detection and tracking of field objects in the known environment of SPL, RoboCup.
    - Particle filter based self-localization module for Nao robot in a known environment.
    - A ROS/Rviz based frontend for Nao robot to achieve user interaction, debugging, robot calibration and visualization of robot state communicated through the network.
    - Interface between NaoQi (Nao Robot software) and V-REP open-source simulator to simulate a Nao robot in V-REP.
- Team Description Paper 2019: `https://tinyurl.com/team-nust-spl-tdp-2019`
- Team Qualification Video 2019: `https://tinyurl.com/team-nust-video-2019`
- Consultancy resource for (UG/PG) students and supervision of local/international summer interns.

ROBOTICS SOFTWARE INTERN   *August 2015 — September 2015*
- Comparison of grid-based and sampling-based search algorithms for path-planning on Nao robots.
- Implementation of a kalman filter based ball tracker for Nao robot.

Exchange Programs & Competitions

Oxford Machine Learning Summer School (OxML) 2023   *University of Oxford + Virtual*
SUMMER SCHOOL STUDENT   *May 2023 — Jul 2023*
- Student of OxML Summer School 2023, Finance + Health track.
- `https://www.oxfordml.school/`

Sakura Exchange Program in Science   *Tohoku University, Sendai, Japan*
RESEARCH EXCHANGE STUDENT   *21 July 2017 — 30 July 2017*
- Visited Tohoku University, Japan, as a robotics research exchange student.
- Attended lectures and tutorials on space robotics, mobile robotics, and micro-electromechanical systems (MEMS).

Standard Platform League (SPL), RoboCup   *Leipzig, Germany*
TEAM MEMBER   *30 June 2016 — 4 July 2017*
- Competed in the international robotics competition RoboCup '16 (held at Leipzig, Germany) as a junior member of Team-NUST '16.
- Team Description Paper: `https://tinyurl.com/tdp-2016`

Skills

Deep Learning	PyTorch / PyTorch Ignite / Explainable AI / Private ML / Document AI / CNNs / Transformers / Diffusion Models / Distributed Computing / SLURM
Frontend/Backend	FastAPI / React / NextJS / S3 / LakeFS / SQL-Alchemy / Supabase / RabbitMQ / PySpark / Docker / Git
LLM Frameworks	LangChain / RAG / vLLM
Edge AI	NVIDIA Deepstream / NVIDIA Jetson / TensorRT
Programming	Python / C++ / C / Javascript / Git
Database	PostgreSQL, MongoDB

Publications

[1] Saifullah, S., Siddiqui, S. A., Agne, S., Dengel, A., Ahmed, S. "Are Deep Models Robust against Real Distortions? A Case Study on Document Image Classification." *26th International Conference on Pattern Recognition (ICPR)*, 2022, pp. 1628–1635. IEEE. Published

[2] Saifullah, S., Agne, S., Dengel, A., Ahmed, S. "Analyzing the Potential of Active Learning for Document Image Classification." *International Journal on Document Analysis and Recognition (IJDAR)*, vol. 26, no. 3, pp. 187–209, 2023. Springer. Published

[3] Saifullah, S., Agne, S., Dengel, A., Ahmed, S. "DocXClassifier: Towards a Robust and Interpretable Deep Neural Network for Document Image Classification." *International Journal on Document Analysis and Recognition (IJDAR)*, pp. 1–27, 2024. Springer. Published

[4] Saifullah, S., Agne, S., Dengel, A., Ahmed, S. "ColDBin: Cold Diffusion for Document Image Binarization." *International Conference on Document Analysis and Recognition*, pp. 207–226, 2023. Springer Nature Switzerland, Cham. Published

[5] Saifullah, S., Agne, S., Dengel, A., Ahmed, S. "PrIeD-KIE: Towards Privacy Preserved Document Key Information Extraction." arXiv preprint arXiv:2310.03777, 2023. Preprint

[6] Saifullah, S., Mercier, D., Lucieri, A., Dengel, A., Ahmed, S. "The Privacy–Explainability Trade-off: Unraveling the Impacts of Differential Privacy and Federated Learning on Attribution Methods." *Frontiers in Artificial Intelligence*, vol. 7, 1236947, 2024. Frontiers Media SA. Published

[7] Saifullah, S., Mercier, D., Agne, S., Dengel, A., Ahmed, S. "Towards Privacy Preserved Document Image Classification: A Comprehensive Benchmark." *International Journal on Document Analysis and Recognition (IJDAR)*, pp. 1–25, 2024. Springer. Published

[8] Agne, S., Dengel, A., Ahmed, S. "The Reality of High Performing Deep Learning Models: A Case Study on Document Image Classification." *IEEE Access*, 2024. IEEE. Published

[9] Saifullah, S., Agne, S., Dengel, A., Ahmed, S. "DocXplain: A Novel Model-Agnostic Explainability Method for Document Image Classification." arXiv preprint arXiv:2407.03830, 2024. Published

[10] Riaz, N., Saifullah, S., Agne, S., Dengel, A., Ahmed, S. "StylusAI: Stylistic Adaptation for Robust German Handwritten Text Generation." *International Conference on Document Analysis and Recognition*, pp. 429–444, 2024. Springer Nature Switzerland, Cham. Published

[11] Hamdani, S. J. H., Saifullah, S., Agne, S., Dengel, A., Ahmed, S. "Latent Diffusion for Guided Document Table Generation." *International Conference on Document Analysis and Recognition*, pp. 368–383, 2024. Springer Nature Switzerland, Cham. Published

[12] Saifullah, S., Agne, S., Dengel, A., Ahmed, S. "DocVCE: Diffusion-based Visual Counterfactual Explanations for Document Image Classification." arXiv preprint arXiv:2508.04233, 2025. Preprint

[13] Saifullah, S., Agne, S., Dengel, A., Ahmed, S. "DP-DocLDM: Differentially Private Document Image Generation using Latent Diffusion Models." arXiv preprint arXiv:2508.04208, 2025. Published

Opensource Contributions/Projects

Projects

ATRIA	https://github.com/saifullah3396/atria	Project
TORCHXAI	https://github.com/saifullah3396/torchxai	Project
DOCXCLASSIFIER	https://github.com/saifullah3396/docxclassifier.git	Project
COLDBIN	https://github.com/saifullah3396/coldbin	Project
DP-DOCLDM	https://github.com/saifullah3396/dpdocldm	Project
ROS/GYM INTEGRATION	https://github.com/saifullah3396/ros_gym.git	Project
MAVROS/MOVEIT INTEGRATION	https://github.com/saifullah3396/mavros_moveit	Project

Contributions

MICROSOFT AIRSIM	https://github.com/microsoft/AirSim/issues/2112	Bug Fix
OCTOMAP CUDA SUPPORT	https://github.com/OctoMap/octomap/pull/257	Feature Contribution
PX4-GAZEBO INTEGRATION	https://github.com/PX4/PX4-SITL_gazebo-classic/pull/362	Feature Contribution
DEEPSHAP EXPLAINER	https://github.com/shap/shap/issues/1479	Contribution