

Yağız Ekrem Dalar

Istanbul, Türkiye – yozydalar2010@gmail.com – ekrem@ethosoft.org

linkedin.com/in/yagizekremdalar – github.com/Yagiz-D-Ekrem-Dolar – ORCID: 0009-0009-3522-4403

Professional Summary

AI developer, researcher, data engineer, and technical project builder working across computer vision, reinforcement learning, natural language processing, multimodal systems, LLM agents, scientific computing, medical AI, and reproducible machine-learning infrastructure. Experience includes designing experimental pipelines, implementing multi-agent architectures, building end-to-end data and model workflows, running large-scale GPU/HPC experiments, preparing academic manuscripts, and translating research prototypes into interactive demos and developer-facing tools. Research activity spans uncertainty-aware reinforcement learning, open-weight LLM safety, Turkish NLP and tokenization, document visual question answering, implicit neural representations, topology-preserving segmentation, federated intrusion detection, vision-language reasoning, computational biology, and agent orchestration.

Research Profile and Current Focus

- **Computer Vision and Multimodal AI:** medical image analysis, semantic and referring-expression segmentation, document understanding, visual-language models, topology-aware learning, implicit neural representations, and synthetic benchmark design.
- **Reinforcement Learning:** temporal credit assignment, ensemble uncertainty, adaptive advantage estimation, PPO, TD3, SAC, DreamerV3, benchmark design, ablation studies, and failure-mode analysis.
- **Natural Language Processing:** Turkish morphology-aware tokenization, sentence-level semantic planning, small-language-model agents, retrieval of tools and skills, transformers, and evaluation of long-form generation.
- **LLM Safety and Agentic Systems:** open-weight tamper-resistance, removal-cost alignment, LoRA/QLoRA experimentation, multi-agent orchestration, context engineering, verifier/fallback systems, and automation pipelines.
- **Scientific and Health AI:** histopathology, cardiomegaly and polyp detection, molecular docking, protein-structure analysis, pharmacological chaperone screening, and research collaboration in health-oriented image/signal processing.
- **Data and Research Engineering:** SQL workflows, data pipelines, experiment tracking, reproducibility manifests, synthetic data generation, CI, structured logging, HDF5, SLURM, CUDA, and TRUBA HPC execution.

Education

Pertevniyal Anadolu Lisesi

High School Education

Istanbul, Türkiye

2024–2029

- Conducting advanced independent and collaborative research in artificial intelligence, machine learning, scientific computing, and software engineering alongside formal secondary education.
- Developing research manuscripts, competition projects, open-source repositories, and reproducible experimental systems.

Research and Professional Experience

Ethosoft / Ethos Research Group

Vice Captain, Data Engineer, and AI Researcher

Istanbul, Türkiye

2025–Present

- Build scalable AI systems, multi-agent architectures, research automation pipelines, and LLM-based decision-support agents.
- Contribute to research projects in reinforcement learning, open-weight LLM safety, Turkish NLP, multimodal document intelligence, computer vision, scientific simulation, and federated learning.
- Prepare datasets, model-training pipelines, evaluation suites, ablation studies, statistical analyses, reproducibility manifests, and technical reports for research submissions.
- Work with GPU/HPC environments, SLURM jobs, containerized execution, experiment logging, and cross-model benchmarking.

Istanbul Technical University - UZDES

Research Collaborator, working with Dr. Yusuf Hüseyin Şahin

Istanbul, Türkiye

Ongoing

- Participate in research-oriented work involving computer vision, image processing, human-centered computing, and machine-learning experimentation.
- Support the transition from research questions to implementable pipelines through dataset preparation, model prototyping, evaluation design, and documentation.

- Collaborate in an academic setting focused on reproducibility, careful benchmarking, and the communication of technical results.

Istanbul Technical University - Health Laboratory

Research Collaborator, working with Prof. Dr. Behçet Uğur Töreyn

Istanbul, Türkiye

Ongoing

- Contribute to health-oriented AI studies connected to image processing, signal processing, machine learning, pattern recognition, and biomedical applications.
- Develop and evaluate computer-vision and deep-learning workflows for medical or health-related data, with attention to preprocessing, validation, model limitations, and reproducibility.
- Bridge software implementation and research reporting by preparing experiments, analyses, and technical documentation.

Uyumsoft

Data Engineer Intern

Istanbul, Türkiye

2025

- Developed SQL-based data systems and supported the optimization of data-processing workflows.
- Worked with structured data, pipeline logic, query design, transformation steps, and operational data-quality considerations.

Neuronauts

Head of Public Relations and AI Developer

Istanbul, Türkiye

2024–2025

- Led community growth and public-facing coordination while contributing to AI projects in natural language processing and computer vision.
- Combined technical development with communication, presentation, project organization, and collaborative team leadership.

Research Manuscripts, Preprints, and Submission Drafts

Uncertainty-Gated Temporal Credit in Reinforcement Learning (UGTC)

2026

Yağız Ekrem Dalar, Feyzi Arda Salihoğlu, Nedim Mutlu Sezer, Ömer Faruk Aksoy, Ahmet Rıfat Öztürk

Introduces a plug-in mechanism that blends short- and long-horizon credit estimates using ensemble disagreement as a state-dependent uncertainty proxy. The work evaluates UGTC with PPO, TD3, SAC, and DreamerV3 across six benchmark suites and reports both successful and unsuccessful cases. The study includes ten-seed experiments, bootstrap confidence intervals, compute-cost analysis, ablations, gate calibration, and a preliminary applicability diagnostic. A workshop-oriented draft was also prepared for the NewInML NeurIPS 2026 context; submission status should be described as a draft unless formally updated.

ForesightLM: Learning Sentence-Boundary Future Energies for Token-Level Autoregressive Generation

2026

Ahmet Rıfat Öztürk, Nedim Mutlu Sezer, Feyzi Arda Salihoğlu, Yağız Ekrem Dalar, Ömer Faruk Aksoy

Presents a lightweight auxiliary objective that trains a token-level causal language model to predict future sentence embeddings at sentence boundaries. Evaluated on TinyStories, WikiText-103, and WritingPrompts-derived data using DistilGPT-2. The paper studies learnability, compatibility with standard language-model loss, semantic reranking, direct future-head-aware reranking, degeneration metrics, bootstrap confidence intervals, and compute accounting. Results show that future semantic prediction is learnable while generation benefits depend on domain and calibration.

Mycelium: Safety Should Be Expensive to Remove - Budget-Dependent Removal-Cost Alignment for Open-Weight LLMs

2026 revision draft

Nedim Mutlu Sezer, Ahmet Rıfat Öztürk, Yağız Ekrem Dalar, Ömer Faruk Aksoy, Feyzi Arda Salihoğlu

Explores removal-cost alignment for open-weight language models: safety should not only appear in behavior, but should be costly to remove while preserving utility. The draft evaluates multiple open model families with static excision, adaptive LoRA repair, safety-erosion sweeps, generation audits, and bypass controls. The work explicitly documents prototype boundaries and negative results, framing the current system as evidence for a design principle rather than a complete deployable defense.

NEDOTURKISHTOKENIZER: Morphology-Aware Hybrid Tokenization for Turkish

submission draft

Anonymous EMNLP 2026

Author list anonymized for double-blind review

Develops a morphology-aware hybrid tokenizer for agglutinative Turkish using a 64K-lemma dictionary, more than 260 suffix patterns, deterministic candidate-based segmentation, and twelve targeted correction layers. The submission reports evaluation on the full TR-MMLU benchmark, five tokenizer baselines, four downstream tasks, a 500-word expert boundary study, and ablation analysis. The reported system improves Turkish-token coverage and reduces unexplained BPE fallback while lowering token fertility.

TR-DocVQA-Synth: A Large-Scale Synthetic Visual Question Answering Dataset for Turkish Business Documents *Anonymous submission draft*

Author list anonymized for double-blind review

Introduces a synthetic Turkish document-VQA benchmark containing 15,000 business-document images and 235,000 question-answer pairs across invoice, contract, quotation, and order-style documents. The research compares Donut, LoRA-tuned PaliGemma-3B, and zero-shot Pix2Struct using Turkish-aware exact match, ANLS, and token-level metrics. The work emphasizes privacy-preserving synthetic generation, document-layout diversity, field-level provenance, and reproducible error analysis.

SPF-Seg: State-Space-Conditioned Potential Field Segmentation with Continuous Boundary Representation and Intrinsic Topology Preservation *2026 journal manuscript*

Author list to be finalized in submission metadata

Proposes a grid-decoder-free segmentation paradigm in which masks are represented as the zero level set of a continuous signed potential field. The architecture combines a selective state-space model, sparse spatial anchors, Fourier-featured implicit neural representation, eikonal regularization, and persistent-homology-based topology loss. The manuscript evaluates thirteen datasets against nine baselines and reports strong boundary and topology performance with a compact parameter count.

Exposing Contextual Amnesia in Vision-Language Segmentation *Double-blind research article draft*

Author information withheld for peer review

Defines and measures “contextual amnesia,” a failure mode in which vision-language segmentation systems recognize isolated attributes but fail on relational, compositional, negated, multi-target, or absent-target prompts. The study introduces a controlled synthetic benchmark with 10,000 images and 52,791 rendered objects, compares six open-source pipelines across five linguistic regimes, evaluates lightweight specialist models, and studies targeted adaptation and late-fusion ensembles.

Robust Federated Intrusion Detection Across Data and Trust Assumptions: A 13,065-Run Paired Benchmark *Double-blind review manuscript*

Author list anonymized for review

Benchmarks eleven federated aggregation rules on UNSW-NB15 and CICIoT2023 under changes in attack type, heterogeneity, malicious-client fraction, participation, identity continuity, scale, and trusted-information assumptions. The archived campaign includes 12,680 main runs and 385 pilot/ablation runs. The paper emphasizes paired seeds, tail-risk reporting, deployability assumptions, held-out attack selection, and the non-portability of method rankings across datasets.

Context-Compiled Small Agents: Efficient Agentic Orchestration with Context Engineering *Research note / manuscript draft, 2026*

Collaborative research draft

Proposes that small language models should be deployed as compiled specialist agents consisting of a model, role contract, context packet, selected tools, selected skills, verifier policy, and fallback plan. The architecture includes a task router, context-packet compiler, tool/skill selectors, schema and policy verification, self-repair, stronger-model fallback, and budget-aware metrics. The paper defines hypotheses and an experimental protocol for comparing large-only, small-only, context-engineered, and verifier-assisted agent systems.

Breast Cancer Histopathology Classification with Convolutional Neural Networks *Technical research report*

Yağız Ekrem Dalar

Documents a CNN-based workflow for benign/malignant histopathology classification, including image resizing, tensor conversion, normalization, bilateral filtering, histogram equalization, grid-distortion augmentation, optimizer comparison, hyperparameter selection, model saving, and evaluation with classification metrics.

Accepted Publications, Workshops, and Research Events

Accepted record as of August 2026. Research outputs, workshop contributions, journal work, technical demonstrations, or related presentation materials have received acceptance from the venues listed below. For double-blind or workshop submissions whose official title-to-venue mapping is not publicly shown in the connected records, the venue name is preserved without expanding or guessing the confidential submission title.

- **UYES Journal:** Acceptance associated with the final UYES version of *Uncertainty-Gated Temporal Credit in Reinforcement Learning (UGTC)*. The work studies state-dependent temporal credit assignment with short- and long-horizon critics, uncertainty gating, multi-backbone evaluation, ablations, reproducibility, and failure analysis.

- **RoboCup:** Accepted robotics-oriented research, workshop, educational-technology, or demonstration contribution connected to the broader autonomous-agent and educational-robotics portfolio.
- **WEROB:** Accepted workshop contribution in the educational-robotics context. Related engineering work includes material-aware project generation, validated microcontroller firmware, code–circuit consistency checking, Wokwi-compatible artifacts, simulation-aware reporting, and high-school-oriented technical explanations.
- **COLM Workshop Program:** Accepted workshop research contribution in language modeling, agentic systems, context engineering, model evaluation, or related machine-learning research.
- **LLA Workshop:** Accepted workshop contribution. The exact official workshop subtitle and title mapping are retained in submission records rather than inferred in this CV.
- **AIMS Workshop:** Accepted workshop contribution associated with the 2026 research portfolio; official naming is preserved in acronym form to avoid expanding the venue incorrectly.
- **RSS 2026 Workshop Program:** Accepted workshop contribution connected to robotics, autonomous systems, simulation, learning, or embodied-intelligence research.

Selected Major Research and AI Projects

BIDIK - AI Agent Ecosystem

Multi-agent AI / automation

Designed as a multi-agent ecosystem for LLM reasoning, task decomposition, decision support, tool use, workflow automation, and end-to-end agent coordination. Work includes agent roles, orchestration logic, structured outputs, verification, and integration with practical software workflows.

UGTC Reinforcement-Learning Framework

Reinforcement learning research

Implemented dual-horizon critics, an ensemble-disagreement gate, EMA-normalized uncertainty, and integrations for multiple RL backbones. Built evaluation pipelines for six benchmark suites, ten random seeds, deterministic evaluation, bootstrap confidence intervals, ablations, throughput measurement, and failure-mode analysis.

ForesightLM

NLP / language generation

Built a sentence-boundary future-semantic objective on top of token-level autoregressive generation. The project includes data preparation for sentence-transition pairs, frozen sentence-encoder targets, contrastive training, candidate generation, semantic reranking, future-head reranking, degeneration metrics, human-evaluation protocol, and reproducibility artifacts.

Mycelium

Open-weight LLM safety

Investigated removal-cost alignment using low-rank interventions, model-family comparisons, static removal tests, utility audits, bounded LoRA repair, safety-erosion sweeps, manual generation review, and explicit bypass controls. The project is documented with strong emphasis on limitations, threat-model boundaries, and reproducibility.

NEDO Turkish Tokenizer

Turkish NLP

Developed a deterministic morphology-aware tokenizer with dictionary-backed roots, suffix patterns, candidate scoring, typed output tokens, canonical morpheme labels, correction layers, intrinsic coverage evaluation, downstream task evaluation, and human boundary annotation.

TR-DocVQA-Synth

Document AI / multimodal learning

Constructed a privacy-preserving synthetic Turkish document-understanding benchmark for invoices, contracts, quotations, and orders. Work covers structured-record generation, HTML/PDF/PNG rendering, question-answer generation, layout diversification, field provenance, model fine-tuning, and metric/error analysis.

SPF-Seg

Computer vision / continuous segmentation

Developed a continuous potential-field approach to segmentation using state-space context mixing, coordinate-conditioned implicit networks, Fourier features, eikonal loss, persistent homology, arbitrary-coordinate queries, and topology-aware evaluation.

RefSeg-Synthetic and Contextual-Amnesia Benchmark

Vision-language reasoning

Created a procedural benchmark for relational, compositional, negated, multi-target, and absent-target referring expressions. The project includes synthetic scene generation, controlled linguistic regimes, large metric panels, zero-shot pipeline evaluation, small-model training, adaptation methods, and specialist ensembles.

CAST / CAST-FL+ Federated Intrusion Detection Benchmark

Federated learning / cybersecurity research

Built an assumption-aware benchmark for robust aggregation methods under poisoning, non-IID data, partial participation, identity reset, client-scale changes, and trusted-root or oracle-count privileges. The project uses paired seeds, cross-dataset ranking analysis, regret and tail-risk metrics, and reproducible run manifests.

Cardiomegaly Detection

Medical computer vision

Deep-learning pipeline for medical-image-based cardiomegaly analysis, including preprocessing, model training, validation, and decision-support-oriented evaluation.

Polyp Detection

Medical computer vision

Real-time computer-vision system for early polyp detection, designed around efficient inference, localization, and clinically relevant error analysis.

Breast Cancer Histopathology CNN

Medical AI

Image-classification study for benign and malignant histopathology images using CNNs, preprocessing and augmentation, optimizer experiments, regularization, and quantitative evaluation.

Protein Analysis AI Agents

Computational biology / agentic AI

Agent-based and reinforcement-learning-oriented workflows for protein analysis, literature-assisted reasoning, candidate prioritization, and biological data interpretation.

G6PD Pharmacological Chaperone and Molecular Docking Demo

Computational biology / drug discovery

Prepared wild-type and R166H mutant G6PD structures, used AlphaFold-oriented structure preparation, screened 3,799 Drug-Central molecules with GNINA, compared docking affinity and selectivity, selected a candidate for analysis, and built a Flask/PyMOL-oriented visualization demo with PDB/SDF parsing and interactive scientific presentation.

D-RIFT Scientific Simulation Infrastructure

HPC / research software

Developed the Phase-0 infrastructure and test skeleton for a high-integrity scientific simulation stack: SI-unit constants, strict Pydantic/YAML configuration, structure-of-arrays particle storage, deterministic shard-invariant random generation, NVIDIA Warp CPU-GPU round trips, validation checks, atomic HDF5 output, JSONL logs, reproducibility manifests, pytest/Ruff CI, SLURM jobs, and evidence-based stage gates.

Solvio - English Learning Application

Educational AI

Adaptive AI-supported language-learning application designed to personalize exercises, feedback, and progression according to learner performance.

Mucilage Prevention System

Environmental AI

AI-based environmental monitoring and prediction project focused on early risk detection, data-driven indicators, and preventive decision support for mucilage-related conditions.

Data and Workflow Automation Systems

Data engineering

SQL-backed pipelines, structured ingestion and transformation, reporting, automation, validation, and integration of AI components with web/API applications.

Personal GitHub Project Portfolio

Primary public GitHub identity: [Yagiz-D-Ekrem-Dolar](#). The account is presented as the personal software and research portfolio. Projects whose repositories are private are described only at the level supplied by the author; projects that are forks or adaptations are explicitly labeled to avoid implying original authorship of upstream work.

DART-RIFT

Python / GPU scientific computing / planetary defense

A high-integrity simulation and Bayesian-inference program focused on recovering probabilistic information about the internal structure of Dimorphos from the NASA DART impact before independent measurements from ESA Hera. The repository implements a deterministic scientific-computing foundation, an SPH shock-physics engine, and material-physics components including strength, porosity, and self-gravity. Its engineering philosophy is based on evidence-gated phases rather than feature completion alone: G0, G1, and G2 are documented through reproducible acceptance criteria, clean repository states, CPU/GPU parity checks, structured evidence reports, and large automated test suites. The project includes SI-unit enforcement, strict versioned configuration, structure-of-arrays particle storage, deterministic and shard-invariant random generation, HDF5 scientific output, JSONL logs, reproducibility manifests, NVIDIA Warp kernels, NumPy FP64 references, SLURM execution, TRUBA GPU validation, conservation tests, analytical shock benchmarks, convergence studies, and explicit model-limit documentation.

Reported engineering work also examines scalability to millions of particles, GPU-memory use, self-gravity cost, numerical truncation error, energy drift, and the distinction between passing a verification gate and making a defensible scientific claim about Dimorphos.

Orbital Drift Lab

JavaScript / browser physics

A compact HTML canvas physics playground for interactively launching luminous bodies under gravity. The application implements gravitational motion, boundary rebounds, particle-to-particle collision behavior, click-to-spawn interaction, motion trails, and live simulation statistics. Although intentionally small, it demonstrates event-driven browser programming, real-time animation loops, numerical state updates, collision handling, user-controlled initial conditions, and visual communication of physical behavior without a heavy framework or server dependency.

WeaveUI Chat

TypeScript / WebSocket / generative interfaces

A multi-user chat environment in which the model does not answer primarily with prose; instead it returns declarative HTML template patches that modify a shared live page for every connected participant. The system defines a lightweight data-slot/data-mode patch protocol for append, prepend, and replace operations, broadcasts model-produced patches through WebSockets, preserves bounded room history and replayable patch state, and supports round-trip interaction through generated forms, quick-reply controls, buttons, and game elements. The client runtime applies partial DOM updates without a virtual-DOM framework, reactivates generated scripts, and protects an external prompt control from being removed by model-generated redesigns. Server-side logic includes provider abstraction, Anthropic/Groq support, multiple owned-key failover for transient limits, fallback wrapping when a model violates the template contract, environment-variable handling, and deployment guidance for persistent WebSocket hosts. The repository also documents the security boundary of running model-produced scripts and treats the project as a local-development and trusted-user research prototype rather than a hardened public service.

UGTC Documentation and Research Site

CSS / research communication / reinforcement learning

Public-facing documentation for *Uncertainty-Gated Temporal Credit*. The underlying research introduces a fast critic using a conservative temporal horizon, a slow critic ensemble using a longer horizon, and an EMA-normalized disagreement gate that selects the appropriate credit scale for each state. The project communicates integrations with PPO, TD3, SAC, and DreamerV3; experiments across Hopper, Ant, Progen, MetaWorld, CarRacing, and Crafter; ten-seed evaluation; bootstrap confidence intervals; throughput and compute accounting; ablations; gate calibration; positive results; negative results; and the boundary condition that uncertainty must vary meaningfully across states to be useful. The repository functions as a research-communication layer for presenting methodology, results, diagrams, and reproducibility information in an accessible web format.

CK3 AI Agent

Game AI / behavior cloning / reinforcement learning

Documentation and data-publication project for an autonomous agent intended to learn the complex mechanics of *Crusader Kings III*. The associated public dataset contains synchronized gameplay video, mouse coordinates and clicks, keyboard events, and temporal alignment suitable for imitation learning and sequence modeling. The development roadmap combines behavior cloning from expert play, reward-driven reinforcement learning, and eventual autonomous campaign execution. The project also covers dataset communication, Hugging Face publication, research documentation, GitHub Pages deployment, and the challenge of converting high-dimensional strategic-game interaction into learnable perception-action trajectories.

Voice Bank App

TypeScript / Next.js / voice interfaces / fintech UI

A modern TypeScript application built with Next.js, React, Tailwind CSS, shadcn-style component tooling, interactive charts, drag-and-drop interfaces, motion components, Three.js/React Three Fiber, and Pipecat voice-client packages. The dependency structure indicates a voice-enabled financial or banking interaction prototype with support for WebRTC and WebSocket voice transport, reusable voice UI components, 3D visualization, globe-style interaction, dashboard charts, date and command interfaces, and production-oriented build/lint scripts. Because the public README is currently empty, the CV limits claims to capabilities visible in the repository configuration rather than asserting unverified product behavior.

New Idea Exchange

JavaScript / corporate innovation platform

A responsive internal-innovation platform that supports idea submission, anonymous feedback, voting, AI-assisted idea analysis, role-based menus, management dashboards, and administration workflows. The application includes a server-side AI-analysis endpoint, protects model credentials by keeping API keys outside the frontend, supports demo-mode fallback when external model access is unavailable, and provides role-specific experiences for employees, managers, and administrators. The concept has also been adapted into private organization-specific demonstrations, creating a reusable product pattern for structured ideation, evaluation, organizational feedback, and innovation-pipeline visibility.

Organization-Specific New Idea Exchange Variants

Private JavaScript repositories

Private variants include the Sabanci, Safi, and Heidelberg Materials editions of the New Idea Exchange platform. They represent organization-specific demonstrations or adaptations of the corporate innovation platform. Since the repositories are private and their implementation is not exposed through the connected source, this CV records the existence and product family without inventing client data, deployment status, confidential integrations, or production use.

A private software project listed in the personal portfolio. Technical details are intentionally omitted because the connected account does not expose the repository content. It is retained to show active private product-development work while respecting confidentiality and avoiding unsupported claims.

Hermes Agent - Personal Fork and Evaluation Workspace*Python / agent systems / forked upstream project*

A personal fork of the Nous Research Hermes Agent project, used as an experimentation and evaluation workspace rather than claimed as an original implementation. The upstream system explores persistent and self-improving agents with tool use, skill creation, cross-session memory, scheduled automation, sub-agent delegation, multiple communication channels, model-provider flexibility, and execution across local, container, SSH, HPC, and serverless environments. Work on the fork supports studying agent architecture, skills, memory, orchestration, deployment, and possible customization while preserving attribution to the upstream authors.

AINavi / FederNet Exploration*Computer vision / planetary navigation / external-project integration*

A public repository containing or tracking a FederNet-based terrain-relative navigation system for lunar and planetary applications. The upstream approach combines convolutional crater detection, crater matching through projective and geometric information, position estimation, and Extended Kalman Filter integration. The repository includes detection, matching, position-estimation, and filtering resources together with lunar imagery and crater-database references. Because the source documentation identifies an external university research origin, this CV presents the repository as an exploration, integration, or study workspace and does not claim original authorship of the FederNet method.

BBVA*Public repository / project placeholder*

A public repository present in the personal account. At the time of review, the connected GitHub metadata exposes no substantive public content or project description. It is therefore listed as a portfolio placeholder without an invented technical scope.

Ethosoft Institutional Project Portfolio

Organization: github.com/ethosoftai. These projects are presented as institutional or team projects associated with Ethosoft. The descriptions distinguish directly documented functionality from repository-level scope inferred only from visible names.

RoboMentor / code2board*Python / educational robotics / LLM-agent toolchain*

A material-aware project-generation system designed for RoboCup Junior and WEROB-style educational robotics. Given an inventory of electronic components and a natural-language request, the system selects a feasible board and project template, creates deterministic pin assignments, generates Arduino C++ firmware, emits Wokwi-compatible circuit files, produces wiring tables and diagrams, runs pin/power/code-circuit/simulation/educational-safety validators, and creates HTML/Markdown reports and student-facing explanations. It can run fully offline through a deterministic heuristic planner while optionally using external LLM providers only to enrich explanations rather than rewrite validated firmware. The project supports multiple microcontroller families, twenty-plus component types, line-following and obstacle-avoidance robots, traffic-light systems, greenhouse monitoring, static simulation fallback, Graphviz/Matplotlib visualizations, metrics logging, Claude-compatible skill packaging, and reproducible workshop-data generation. It explicitly records when a real electrical simulation has not occurred and documents hardware-safety limitations for educational use.

DataMedX2 Agent*Clinical agent orchestration / medical decision-support prototype*

A hackathon-oriented clinical orchestration and physician-dashboard application. A doctor can enter a free-form request containing different patient-code formats; the system normalizes the identifier, retrieves the patient record, runs specialized clinical sub-agents, and assembles an evidence-linked output. The agent network covers patient retrieval, clinical summarization, laboratory review, treatment and medication context, oncology status, risk triage, timeline construction, report generation, and safety checking. The interface includes source snippets, red/yellow/green risk cards, a longitudinal patient journey, Markdown/PDF reporting, and demo metrics such as evidence count, risk count, timeline events, and reporting efficiency. The project treats extreme laboratory values as signals requiring clinician verification rather than autonomous diagnoses and separates runtime secrets from the repository through environment configuration.

BIDIK Ecosystem, BIDIKReklam, and bidik_calendar*Multi-agent automation / communication / scheduling*

An institutional project family centered on the BIDIK AI-agent ecosystem. The broader system combines LLM reasoning, task decomposition, automation, structured tool use, decision support, and end-to-end orchestration. BIDIKReklam represents an advertising, outreach, campaign, or public-communication component, while bidik_calendar provides a calendar and scheduling-oriented component. Together, the family demonstrates how a core agent platform can be divided into user-facing communication, planning, scheduling, and workflow modules. Exact production integrations are not asserted where public documentation is unavailable.

LunaLocater

Python / lunar navigation / computer vision

A FederNet-oriented lunar localization pipeline intended to estimate position from visual observations in environments where GPS-like infrastructure is unavailable. The documented architecture analyzes lunar imagery, detects and marks craters using YOLO-style object detection, compares observed crater geometry with satellite or catalog data, and combines matching outputs with optimization and an Extended Kalman Filter to estimate coordinates. The repository is organized around notebooks for settings, dataset creation, crater catalogs, optional model training, matching, position finding, Kalman filtering, end-to-end execution, and test automation. The project connects computer vision, geometric matching, state estimation, lunar datasets, and autonomous navigation for future robotic or spacecraft missions.

UGTC - Ethosoft Research Repository

Reinforcement learning / uncertainty-aware optimization

Institutional repository associated with the UGTC research program. It supports documentation, code organization, reproducibility, and communication of adaptive temporal credit assignment. The research infrastructure covers dual-horizon value estimation, ensemble disagreement, state-dependent gating, multiple reinforcement-learning backbones, benchmark automation, random-seed control, statistical intervals, ablation studies, compute reporting, failure analysis, and applicability diagnostics.

SA-GAN

Generative modeling research

Institutional generative-adversarial-network experiment. The repository name indicates self-attention or spatial-attention GAN research; because a public README was not available in the connected source during review, the CV records it as a generative-modeling project without inventing architecture, dataset, metric, or deployment claims.

Astro

Space and astronomy software prototype

Institutional repository for an astronomy or space-oriented project. It complements the team's work in lunar localization, planetary navigation, simulation, and AI for scientific domains. Public metadata confirms the repository but does not provide a sufficiently detailed README, so the description remains at verified project-family level.

3D-Game

Interactive 3D software project

A sizable institutional repository for an interactive 3D game or simulation prototype. It represents experience with real-time interaction, scene design, gameplay or simulation logic, and visual software development. Specific engine, gameplay mechanics, and release state are not inferred without repository documentation.

karatahta-frontend

Frontend engineering / collaborative interface

Frontend repository for the Karatahta project. It demonstrates institutional work on user-facing web architecture, interface components, product presentation, and integration with a broader application backend. Public repository metadata does not expose a complete product specification, so confidential or undocumented functionality is not claimed.

ethosoft-web

Institutional web platform

The public web repository for Ethosoft's organizational presence, project communication, and institutional identity. The project supports presenting research, software, team activity, and public-facing material through a maintained web application.

Harvard-Hackathon

Hackathon prototype

An institutional hackathon repository used for rapid prototyping, demonstration, collaborative implementation, and time-constrained product development. The repository title is preserved as published; no competition result or affiliation is inferred solely from the name.

sinai and ab-ihayat

Institutional experimental repositories

Additional public Ethosoft repositories representing active or archived institutional experiments. Since the connected GitHub source does not expose a reliable public project description for these repositories, they are included as named portfolio items without fabricated functionality.

Technical Skills

Programming and Query Languages: Python, SQL; experience integrating command-line tools, APIs, scripts, notebooks, web backends, Android-oriented applications, and configuration-driven systems.

Core Machine Learning: supervised learning, deep learning, transfer learning, self-supervised and auxiliary objectives, ensemble learning, uncertainty proxies, hyperparameter optimization, class imbalance handling, calibration, ablation studies, cross-validation, error analysis, and benchmark design.

Computer Vision: image classification, object detection, semantic segmentation, referring-expression segmentation, medical imaging, image preprocessing, synthetic scene generation, OCR-free document understanding, visual-language models, implicit

neural representations, Fourier features, state-space vision models, topology-aware losses, level sets, persistent homology, and real-time inference.

Reinforcement Learning: PPO, TD3, SAC, DreamerV3, GAE, temporal credit assignment, short/long-horizon critics, ensemble disagreement, adaptive gating, value estimation, continuous control, procedural environments, imagined trajectories, seed-based evaluation, and reward-density/failure diagnostics.

Natural Language Processing:

transformers, tokenization, Turkish morphology, subword analysis, BPE, WordPiece, SentencePiece concepts, sentence embeddings, contrastive learning, long-form generation diagnostics, semantic reranking, NER, sentiment classification, NLI, language-model evaluation, and multilingual document understanding.

LLMs and Agents: LLM agents, multi-agent systems, role contracts, task routing, context packets, tool/function calling, skill retrieval, progressive disclosure, verifier/fallback loops, structured outputs, LangChain-style orchestration, local models, small language models, LoRA/QLoRA, model evaluation, and open-weight safety research.

Data Engineering: SQL development, ETL/ELT concepts, data cleaning, dataset versioning, train/validation/test splitting, leakage controls, structured schemas, Parquet/CSV/JSON/JSONL, synthetic data generation, data validation, pipeline automation, HDF5, and reproducibility manifests.

Scientific Computing and HPC: NumPy vectorization, Monte Carlo simulation, deterministic RNG, GPU/CPU execution, CUDA-oriented workflows, NVIDIA Warp, SLURM, TRUBA HPC, Apptainer containers, node-local temporary storage, job evidence publishing, compute-cost accounting, and batch experiment management.

Medical and Computational Biology: histopathology classification, cardiomegaly and polyp detection, protein structure preparation, AlphaFold-oriented workflows, GNINA molecular docking, ligand screening, affinity/selectivity analysis, PDB/SDF processing, Biopython, PyMOL visualization, and pharmacological-chaperone exploration.

Frameworks and Libraries: PyTorch, TensorFlow, Keras, scikit-learn, OpenCV, Hugging Face Transformers, sentence-transformers, LangChain, FastAPI, Flask, Streamlit, XGBoost, LightGBM, CatBoost, Dask, Ray, Optuna, MLflow, Weights & Biases, NumPy, Pandas, Matplotlib, Seaborn, Pydantic, pytest, Ruff, Biopython, PyVista, and Three.js/WebXR-oriented export workflows.

Research Methods: literature review, research-question formulation, controlled synthetic benchmarks, paired experimental design, bootstrap confidence intervals, Wilcoxon/Friedman-style comparison concepts, rank and regret analysis, tail-risk reporting, model-card-style limitations, threat models, reproducibility checklists, compute budgets, human-evaluation protocols, and academic manuscript preparation in LaTeX.

Software Engineering: modular architecture, CLI design, API integration, schema validation, structured logging, CI pipelines, test automation, reproducible environments, documentation, Git/GitHub, issue-oriented development, web applications, developer tooling, and report generation.

Research Infrastructure and Experimental Practice

- Design and execution of large experiment matrices with multiple algorithms, datasets, attack conditions, client scales, random seeds, and ablations.
- GPU/CPU training and evaluation using TRUBA HPC, SLURM jobs, containerized environments, deterministic configurations, and workload-specific manifests.
- Statistical reporting with means, standard deviations, bootstrap confidence intervals, paired comparisons, rank analysis, regret analysis, and explicit negative-result reporting.
- Reproducibility through fixed seeds, hashed inputs, versioned configurations, structured logs, environment records, checkpoint management, and separation of pilot, ablation, and final benchmark evidence.
- Technical communication through LaTeX manuscripts, engineering notebooks, pitch decks, posters, project reports, README files, diagrams, and presentation-oriented demos.

Awards and Distinctions

- **TEKNOFEST 2025:** Third Place and Best Presentation recognition.
- **EU Code Week Türkiye:** First Place.
- **EU Code Week Europe:** First Place.
- **Living Lab:** Second Place.
- **Mindbricks:** First Place.
- **DataMedX:** First Place.
- **ITU MostrAI:** Jury Award.
- Multiple national and international hackathon, innovation, and AI-project recognitions.

Leadership, Communication, and Collaboration

- Team leadership as Vice Captain at Ethosoft/Ethos and former Head of Public Relations at Neuronauts.
- Experience coordinating multidisciplinary work across AI research, data engineering, software development, scientific writing, and presentation preparation.
- Preparation of technical presentations, posters, competition materials, pitch decks, engineering notebooks, manuscript drafts, and project demonstrations.
- Ability to translate complex research ideas into implementable milestones, reproducible experiments, and audience-appropriate explanations.

Languages

Turkish	Native
English	Professional working proficiency
Romanian	Limited proficiency
German	Limited proficiency

Professional Links and Research Identity

Email	yozydalar2010@gmail.com
Professional Email	ekrem@ethosoft.org
LinkedIn	linkedin.com/in/yagizekremdalar
GitHub	github.com/Yagiz-D-Ekrem-Dolar
Ethosoft GitHub	github.com/ethosoftai
ORCID	0009-0009-3522-4403