ORIGINAL RESEARCH **Open access**

K-EGoT: Knowledge-Grounded Thought Alignment for Automated Safety Modeling of High-Altitude Solar Drones

Antonis Diamantis[1*], Katerina Pappa[1], Nikolaos Christou[2], Dimitris Papadopoulos[3]

Abstract

With the widening operational scope of high-altitude solar drones, upholding their safety has become an absolute priority. Conventional safety modeling, which relies on manual expert scrutiny, cannot keep pace with accelerating design timelines. Although Large Language Models (LLMs) offer automation, state-of-the-art reasoning architectures such as the Graph of Thoughts (GoT) are overly broad and lack the specialized disciplinary insight needed for meaningful deployment. To remedy this shortfall, we propose K-EGoT, a framework that ties LLM inference to an auditable, domain-anchored corpus of knowledge. Our technique establishes a "Safety Rationale"—a mandatory, traceable bridge connecting LLM-generated model extensions with expert-vetted safety precepts. We subsequently refine a bespoke model using an original "thought process alignment" methodology, applying Direct Preference Optimization (DPO) to assess the soundness of these rationales and ensure the model's inferential chain mirrors expert judgment. Evaluated on a high-fidelity dataset for the flight control–energy coupling challenge, our 7B K-EGoT model registered a Safety Extension Score (SES) of 92.7, substantially eclipsing the 84.7 mark attained by vanilla GoT prompting. Our study delivers a trustworthy, inspectable solution for automated safety modeling for this vital class of aerial platforms.

Keywords High-altitude solar drones, Flight control-energy coupling, Safety modeling, Safety analysis

*Correspondence:
Antonis Diamantis
antonis.diamantis@aua.gr
[1] Department of AI for Cotton and Tobacco Leaf Quality Grading, School of Agriculture, National and Kapodistrian University of Athens, Athens, Greece
[2] Department of Smart Tillage and Soil Compaction Data Science, Faculty of Agriculture, Aristotle University of Thessaloniki, Thessaloniki, Greece
[3] Department of Precision Harvest Monitoring, Faculty of Agriculture, University of Patras, Patras, Greece

Introduction

High-altitude solar drones, often designated as High-Altitude Long-Endurance (HALE) vehicles, are platforms engineered for protracted missions in the stratosphere. As depicted in **Figure 1**, their most conspicuous physical trait is an exceedingly high aspect ratio, with the wings structurally elongated and slender to accommodate expansive photovoltaic solar panels. This configuration permits diurnal flight and energy replenishment via solar radiation, while depending on stored battery power to sustain flight throughout the night. As the use of these platforms proliferates, the demand for uncompromised safety verification becomes paramount. Unlike standard unmanned aircraft, high-altitude solar drones must continually account for energy limitations across their operational envelope. Hence, the foremost safety hazard is not confined to discrete elements but originates from the deep-seated, intricate dynamic interdependency between the Flight Control System (FCS) and the Energy Management System (EMS)—the flight control–energy coupling. When system defects escape detection during early architectural definition, they can trigger disastrous malfunctions, with rectification expenses eclipsing initial prevention costs by orders of magnitude. While conventional safety assessment

This is an open access article distributed under the terms of the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License (CC BY-NC-SA 4.0), which permits non-commercial use, distribution, reproduction, and adaptation in any medium or format, provided that appropriate credit is given to the original author(s) and source, a link to the license is provided, and any changes made are indicated. Adapted material must be distributed under the same license.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

protocols, such as Fault Tree Analysis (FTA) [1] and Failure Mode and Effects Analysis (FMEA) [2], persist, their orthodox execution suffers from an inherent drawback. These instruments are habitually applied as post-hoc design validation exercises, disconnected from the formative stages of system architecture. This procedural chasm implies that safety imperatives are not woven into the design framework from the outset, often delaying the identification of design-embedded flaws that require expensive, disruptive rework. To address this deficiency, contemporary systems engineering champions the integration of Safety Modeling at the outset of the design lifecycle. Specifically, SysML-grounded safety modeling provides a regimented, model-centric framework for conducting an all-encompassing, holistic system dissection. By systematically codifying safety stipulations and boundaries within the system model itself, this tactic aspires to embed reliability by design rather than retroactively inspecting for it, thereby fortifying dependable future functioning. Yet, the tangible enactment of this paradigm persists as a formidable obstacle, given that the work of discerning relevant dangers and rendering them into the rigorous semantics of a SysML model endures as a resource-draining pursuit, profoundly contingent on the seasoned intuition of veteran safety practitioners and exacting a heavy temporal toll.

The rise of Large Language Models (LLMs) heralds a promising avenue for mechanizing this complex modeling operation. Nevertheless, superior modeling fidelity cannot be secured through off-the-shelf models or elementary prompt crafting. Sophisticated inferential frameworks such as Graph of Thoughts (GoT) [3] remain too unspecialized, bereft of the domain-specific underpinning indispensable for this refined safety arena. Analogously, refinement procedures like Direct Preference Optimization (DPO), when aimed exclusively at the terminal output (a tactic termed “behavioral alignment”), might yield a formally correct product while bypassing the integrity of the underlying logical sequence. For high-altitude solar drones, where an AI hazarding a fortuitous correct answer is wholly unacceptable, this restriction is pivotal. Our investigation builds atop these robust scaffolds, yet our distinct offering is a foundational reorientation from “behavioral mimicry” to “thought process alignment.”

In contrast to orthodox strategies that often merely adjust model weights on curated expert exemplars to replicate accurate answers, our Knowledge-Enhanced Graph of Thoughts (K-EGoT) construct directly confronts the opaque “black box” conundrum. We institute the “Safety Rationale”—a verifiable, reconstructable inferential waypoint that renders the model’s logic both perspicuous and entirely examinable. The pivotal breakthrough lies in our application of this mechanism: we recalibrate the DPO’s optimization objective away from the final deliverable (conduct) and steer it squarely toward the caliber of the inferential trajectory itself, coercing the model to genuinely assimilate profound expert heuristics rather than shallow, misleading correlations.

© 2023 The Author(s). | Published by Agro Intelligence Research Press



Figure 1. Typical high-altitude solar drones. reproduced from Ref. [4].

We undertook a rigorous empirical assessment of a purpose-built corpus focused on genuine operational situations involving high-altitude solar drones. The outcomes demonstrate that our K-EGoT architecture, which leverages fluid exploration to disentangle convoluted, coupled issues, markedly outperforms both static and standard fine-tuning regimes. This finding robustly illustrates that our “thought process alignment” embodies a more decisive and efficacious synthesis of disciplinary insight than mere behavioral alignment or isolated general reasoning faculties.

The main contributions of this paper are as follows:

• We introduce a technique to augment the Graph of Thoughts (GoT) architecture by inserting a “Safety Rationale,” which facilitates the profound interweaving of generic LLM inferencing with a verifiable, domain-dedicated drone safety knowledge repository.

• We substantiate that deploying the “Safety Rationale” as the substrate for Direct Preference Optimization (DPO) enables harmonizing the model’s inferential sequence with specialist logic, constituting a sturdier regimen than standard behavioral alignment.

• We compile and disseminate a premium-grade evaluation benchmark for the flight control–energy coupling safety dilemma, encompassing a domain knowledge corpus that supplies a worthwhile asset for ensuing investigations. The benchmark’s trustworthiness is underpinned by a substantial inter-rater agreement coefficient (Fleiss’ Kappa = 0.82).

• Via exhaustive empirical trials, we furnish concrete validation that our knowledge-infusion methodology on a 7B model outstrips conventional fine-tuning and prompting reference points on an identical foundation model, accentuating the primacy of verifiable inference over generalized capability within the sphere of drone safety modeling.

The paper is organized as follows:

© 2023 The Author(s). | Published by Agro Intelligence Research Press

Introduction: The cardinal importance of safety amid the expanding deployment of high-altitude solar UAVs is outlined. It is underscored that customary safety modeling reliant on manual expert dissection struggles to synchronize with brisk development cadences, while extant large language model (LLM) reasoning architectures (exemplified by Graph of Thoughts, GoT) prove excessively generic owing to an absence of sector-specific erudition; The ambition of this inquiry is expressly articulated as advancing a K-EGoT construct to bridge these identified lacunae, coupled with a succinct outline of the paper's subsequent structural logic.

Background and related work: recapitulates prevailing safety dissection instruments (e.g., Fault Tree Analysis, FTA; Failure Mode and Effect Analysis, FMEA) primarily enlisted for post-design corroboration, detached from early system architectural constraints, and traces the genealogy of LLM inferential approaches from Chain of Thought (CoT) through Tree of Thoughts (ToT) to GoT. Concurrently, it identifies shortcomings in the current LLM deployment when confronting structured representations such as SysML state machine diagrams, and the deficits in verifiability and traceability required for UAV airworthiness endorsement, thereby establishing the case for introducing the K-EGoT framework.

Intelligent safety modeling method: the nucleus of the K-EGoT construct is expounded with precision, embracing two principal phases—"domain expert model cultivation" (elevating a foundational LLM into a domain authority through supervised fine-tuning and direct preference optimization) and "dynamic safety extension inference" (performing safety elaboration tasks with the cultivated expert model serving as the linchpin); The systematic formulation of the UAV safety knowledge repository is elucidated (origin extraction, taxonomic categorization and structuring, deduplication and conflict reconciliation, formal encoding), and the concrete mechanics of the framework's training and inference conduit are unpacked, including the assembly of hybrid-length reasoning chain corpora and the two-tier refinement anchored in the "Safety Rationale."

Experiments and analysis: experimental configurations are engineered to corroborate the potency of the K-EGoT construct. Initially, the experimental setup is detailed (assembly of a high-reliability corpus featuring expert annotation, attaining a Fleiss' Kappa of 0.82; formulation of two reference categories: lexical prompting and fine-tuned; grounded in Qwen2-7B-Instruct); Thereafter, investigative queries are resolved across three experimental modalities: comparative performance assays confirming that K-EGoT (SES 92.7) decisively surpasses baseline configurations; Ablation studies quantifying the respective contributions of GoT inference, rationale alignment, and dynamic reasoning; qualitative case dissection confirming the caliber of safety augmentations while noting constraints about conflict resolution criteria.

Discussion: deliberates on the insights gleaned, stressing the demonstrated potency of fusing a generalized inferential engine with a domain-anchored knowledge corpus for specialized sectors, along with the decisive import of authenticating the inferential chain (rather than merely the conclusion) for UAV safety certification; Scrutinizes threats to validity (construct validity is fortified by expert consensus, internal validity is regulated through a uniform model and parameter set, and external validity is tackled by outlining a roadmap for cross-domain transference); Simultaneously, the tangible benefit of K-EGoT in mechanizing front-end safety dissection, curtailing overhead, and truncating the design cycle is explicated.

Conclusion: encapsulates the paper's central contribution—the articulation of the K-EGoT framework to surmount the challenge of automated safety modeling for high-altitude solar UAVs, alongside the empirical corroboration of its edge; Recognizes present constraints (e.g., complications in navigating implicit precedence conflict rubrics); Delineates prospective investigative trajectories (engineering a priority arbitration apparatus, amalgamation with formal verification engines); Finally, appends particulars regarding author roles, funding sources, data accessibility.

Background and Related Work

Background

© 2023 The Author(s). | Published by Agro Intelligence Research Press

Safeguarding high-altitude solar drones is a mission of utmost criticality, given their use in protracted stratospheric operations spanning weeks or even months. These aerial vehicles contend with severe environmental swings and lean on sophisticated autonomous decision-making without persistent human guidance. Hence, realizing endogenous safety—weaving dependability into the very fabric of the system from inception—stands as the preeminent obstacle to securing their mission success. To fulfill this ambition, a regimented and ceremonialized methodology is essential, positioning the Systems Modeling Language (SysML) as a pivotal instrument [5]. Serving as the prevailing benchmark for model-centric engineering in complex aerospace systems, SysML provides an extensive set of graphical notations to rigorously delineate, analyze, and document the intricate dynamics of these drones, preserving design coherence and streamlining early-stage validation. Yet harnessing SysML to capture the safety profile of high-altitude solar drones effectively poses a vexing challenge. The cardinal snag lies not in the formalism itself, but in ceremonially encoding the system's most acute and convoluted peril: the entrenched flight-control–energy coupling. The navigational directives of the Flight Control System (FCS) are inexorably bound to the live status of the Energy Management System (EMS), encompassing variables such as the battery's State of Charge (SOC) and the on-hand solar influx. In turn, the EMS's prowess in scavenging and stewarding energy is directly modulated by the aircraft's pose and trajectory, parameters dictated by the FCS. Attempting to manually wield SysML apparatuses, most notably the State Machine Diagram, to exhaustively catalog and depict every hazardous interaction within this densely coupled decision space is a Herculean endeavor, replete with the risk of overlooking veiled safety preconditions.

Related work

Traditional and model-based safety analysis

Ceremonial safety vetting of avionic systems is a well-established domain guided by exacting protocols and techniques. Bedrock practices encompass Functional Hazard Analysis (FHA) [6], Fault Tree Analysis, and Failure Mode and Effect Analysis. Prescribed by cardinal aerospace safety edicts such as ARP-4761 [7], these instruments are indispensable for identifying and neutralizing threats. To heal the disjuncture in which these methods stand apart from design schemata, the arena of model-based safety analysis has emerged, seeking to integrate safety dimensions directly into design models via mechanisms such as SysML Profiles [8]. Scholars have fruitfully leveraged profiles to embed safety intelligence into models [9, 10] and to buttress hazard analysis [11].

Limitations: Notwithstanding that these model-oriented techniques erect a semantic tether between safety artifacts and design components, the connection lacks generativity. The safety specialist must yet manually decode standards, isolate germane constraints, and render them into the model utilizing the profile's stereotypes and markers. The scaffolding neither proactively advances safety enhancements nor deliberates on the exhaustiveness of the appraisal, remaining a passive archival contrivance persistently yoked to strenuous manual expert labor.

Defect detection of unmanned aerial vehicles

Anomaly flagging in UAV flight data grapples with a triad of impediments: the heterogeneity of irregularity typologies limits a single model's purview, the scarcity of annotated samples curtails the effectiveness of supervised learning, and transducer disturbances corrupt model fidelity. Incumbent scholarship has amplified detection breadth through blended architectures and multi-model ensembles. By way of illustration, the CA-1DCL-EF framework notched a precision of 94.22% when pinpointing bias and drift irregularities [12, 13]. Harnessing unsupervised and semi-supervised learning to reduce reliance on tagged exemplars, the STARE framework achieved an F1 score of 89.34% on the ALFA collection. In contrast, the VAE construct achieved an accuracy of 92.31% on imbalanced datasets [14, 15]. In parallel, data refinement and descriptor engineering are incorporated to fortify disturbance immunity. The marriage of Savitzky-Golay filtration with DCNN-AE yields an F1 score of 99.89%, and MIC-informed descriptor curation reduces the spurious alarm rate by 8% [16, 17].

Limitations: The foregoing inquiries focus on anomaly identification and mitigation throughout the operational lifecycle of unmanned aerial vehicles but fail to embed safety reflection into the system's architectural and design phases. This ethos

© 2023 The Author(s). | Published by Agro Intelligence Research Press

is fundamentally anchored in algorithmic detection of deviant signatures that have surfaced or are brewing within transducer data streams, without the capacity to preemptively obviate or forestall intrinsic risks birthed by platform design blemishes, erroneous control system idealizations, or ill-advised component picks. Accordingly, wedding “design safety” with “operational safety” and ensuring sturdiness and malfunction tolerance via measures such as safety audits at the modeling stage embodies a decisive approach to amplifying the aggregate safety of unmanned aerial vehicle ensembles at the root, and equally signifies a conspicuous lacuna in the extant investigative corpus.

LLM applications in software engineering: Testing and safety

The dawn of Large Language Models (LLMs) [18] has revealed uncharted corridors for automating erudition-intensive tasks in software engineering, encompassing software analysis [19, 20] and the distillation of safety requirements from specifications [21].

Limitations: Their latent capacity notwithstanding, these deployments habitually operate on natural-language briefs or raw code. They remain intrinsically unequipped to grapple with the regimented, diagrammatic grammar of SysML state machine depictions. Beyond this, they commonly fall short of the exacting, reconstructable lineage mandated for drone platform accreditation, as their inferential trajectory stays shrouded and lacks explicit mooring in a vetted safety canon or domain-anchored knowledge repository.

Evolution and limitations of LLM reasoning paradigms

To elevate accomplishment on complex undertakings, LLM inferential strategies have evolved from rectilinear Chain of Thought (CoT) [22] toward more sinuous, multi-forked structures such as Tree of Thoughts (ToT) [23] and Graph of Thoughts (GoT). These scaffolds amplify generic conundrum resolution by sanctioning the charting, weighing, and coalescence of sundry deduction corridors. Constructs such as self-refinement [24] and self-evaluation [25] are habitually enlisted within these edifices to burnish midway artifacts.

Limitations: These vigorous archetypes persist as all-purpose armatures. Their introspective appraisal competences lean upon bland heuristics plucked from the model’s pretrained corpus, which proves woefully scant for terrains exacting profound specialization, as epitomized by drone safety. Irrespective of its branching inquiry, ToT’s heuristic-steered junction vetting can rashly jettison globally preferable conduits when domain acumen is non-existent. To illustrate, within a power stewardship setting, a superficially “wasteful” gambit, such as a prophylactic altitude drawdown, might embody the most judicious long-term maneuver, a nuance certain to elude a one-size-fits-all heuristic. This lays bare a well-delineated investigative chasm: the imperative for a contrivance to anchor these lithe inferential armatures in a confirmable, domain-bespoke bedrock of expertise.

Intelligent Safety Modeling Method

Overall architecture: The Knowledge-Enhanced Graph of Thoughts (K-EGoT) framework

To master the trial of ceremonially and overtly infusing tacit safety design canons into nascent-stage design schematics for high-altitude solar drones, this investigation puts forward a trailblazing Knowledge-Enhanced Graph of Thoughts (K-EGoT) armature. As shown in **Figure 2**, this armature aims to seamlessly meld the formidable generative-inferential muscle of Large Language Models (LLMs) with the abstruse specialist lore of the drone safety sphere, thereby enabling mechanized, reliable safety amplification of SysML state machine diagrams.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

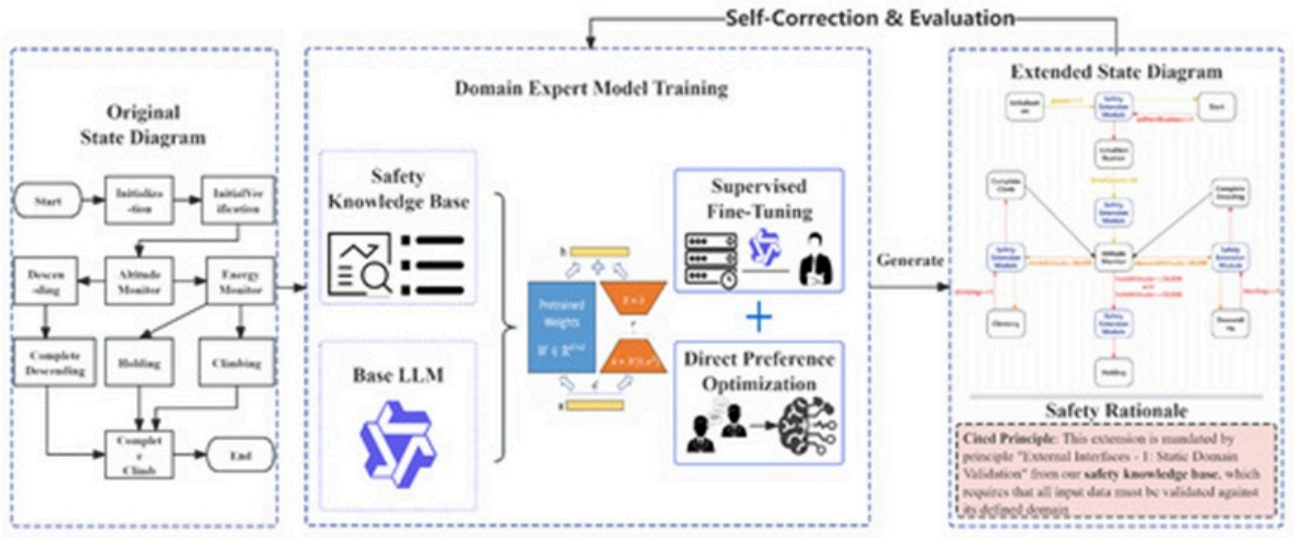

Figure 2. Overall Architecture of the K-EGoT framework, illustrating the two-phase process: (1) training a domain-expert model by aligning it with a safety knowledge base, and (2) using the expert model in a dynamic reasoning loop to extend an initial SysML diagram.

The actualization of this armature engulfs two nuclear stages, pellucidly mirroring the design dictum of “first school the expert, then commission the expert.”

In Phase 1 (domain expert model training), we aim to elevate a universal Base Large Language Model (Base LLM) into a disciplinary savant steeped in the safety of high-altitude solar drone flight-control–energy-coupling. This stage ignites with a pair of cardinal inputs: our curated Drone Safety Knowledge Base and a Base LLM (e.g., Qwen2-7B). We thereupon espouse a Two-Phase Optimization Strategy: first, we execute rudimentary capability harmonization via domain corpora through Supervised Fine-Tuning (SFT); subsequently, via Direct Preference Optimization (DPO), gauged by the soundness of the “Safety Rationale,” we lock in a robust resonance with specialist inferential choreographies. The culminating output of this stage is a thoroughly domain-seasoned Domain Expert Model.

In Phase 2 (dynamic safety extension reasoning), we field the seasoned Domain Expert Model as the hub “execution engine” to shoulder concrete safety expansion missions ensconced within a regimented graph blessed with kinetic self-amendment and sapient probing aptitudes. This stage ingests an Initial SysML State Machine Diagram as its feedstock. Our execution engine ignites its routine, featuring a quintessential internal cycle that oscillates between crafting an expansion blueprint and a safety rationale, followed by self-amendment and appraisal tied to the knowledge repository. This loop churns through Dynamic Exploration and Optimization until a superlative resolution surfaces. In its denouement, this stage dispenses a brace of vital outputs: an enriched safety state machine diagram suffused with endogenous safety implications, and a matched Traceable Safety Rationale.

Knowledge base curation

The cornerstone of our knowledge-enhancement strategy lies in an auditable, discipline-specific knowledge repository. The fidelity and architecture of this corpus are decisive for anchoring the LLM’s inferential operations. We devised a methodical pipeline to assemble this knowledge base, converting broad-spectrum safety directives into a regimented configuration comprising codified objects endowed with unique identifiers and prescribed attribute fields.

The assembly pipeline traversed four cardinal stages:

© 2023 The Author(s). | Published by Agro Intelligence Research Press

· Source Identification and Extraction: The undertaking began with identifying authoritative sources, including umbrella aerospace safety codes (e.g., ARP-4761) and proprietary design dossiers for high-altitude solar drones. A triumvirate of domain authorities exhaustively scrutinized these materials to extract unambiguous safety maxims, requirements, and failure-mode narratives germane to software engineering.

· Classification and Structuring: The harvested raw maxims were then sorted into a regimented taxonomy. We established a pentad classification scheme aligned with vital dimensions of the system's constitution: (1) External Interfaces, (2) Functional Logic, (3) Functional Hierarchy, (4) States and Modes, and (5) Flight Control–Energy Coupling. Individual maxims were allocated to one or several groupings.

· De-duplication and Conflict Resolution: Maxims gleaned from heterogeneous origins routinely exhibited semantic intersections. Experts jointly sifted through the categorized maxims to collapse redundancies and rephrase them into a normalized lexicon. Discernible antagonisms between maxims (e.g., a safety maxim stipulating deactivation of a subsystem contra a functional maxim insisting on its continued operation) were reconciled through the articulation of overt precedence doctrines and context-bound applicability caveats.

· Formalization: As the concluding step, each refined maxim was recast into a machine-interpretable entity carrying a distinct identifier, a compressed natural-language articulation, and disciplined fields (e.g., Hazard, Context, Mitigation Guideline). This formal encoding proves vital for empowering the mechanized substantiation of the KBGrounding yardstick, as examined in Section 4.2.

This painstaking assembly sequence certifies that our knowledge repository is far more than a textual miscellany; it serves instead as a disciplined and verifiable underpinning for our K-EGoT framework. To underwrite a systematic failure mode dissection, the knowledge base is organized into five nuclear groupings: External Interfaces, Functional Logic, Functional Hierarchy, States and Modes, and the decisive Flight Control–Energy Coupling. Each grouping houses bespoke tabulations of evaluative benchmarks, with certain sections exhibiting finer-grained partitioning. By way of example, the "Functional Logic" grouping is further decomposed into four viewpoints, among them control computation and processing rationale, whereas "Functional Hierarchy" is scrutinized in light of series and parallel associations. This stratified construction thereby forms an exhaustive and disciplined knowledge edifice intended to pilot and corroborate the AI model's inferential flow. The meticulous safety benchmarks developed through this methodology are presented in Appendices A and B. These benchmarks form the core of our knowledge repository, providing verifiable "factual anchorage" for downstream intelligent inference.

K-EGoT training and reasoning pipeline

With a domain-knowledge substrate firmly in place, the K-EGoT framework "brings to life" this knowledge through an elaborate training and inference conduit, empowering it to guide and polish the safety modeling enterprise adaptively.

Phase 1: Base model training based on thought process alignment

The opening phase of the K-EGoT framework involves coaching a base model to internalize an "expert-echelon" safety mentality. This stage aims to ensure the model moves beyond mere replication of accurate outputs, genuinely comprehending and regenerating the ingrained logic that guides safety professionals in their adjudicative processes.

Dataset Construction for High-Altitude Solar Drone Safety Scenarios: To equip the LLM with command over the distinctive safety reasoning required by high-altitude solar drones, we assembled an augmented state-machine diagram corpus. This corpus integrates both protracted and condensed Chains of Thought (CoT), which jointly supply the elemental inferential pathways for the subsequent "Graph of Thoughts" probing.

The long-chain corpus was synthesized in the initial step. We leveraged the Deepseek-R1 model, provisioning it with bespoke scenarios for high-altitude solar drones (e.g., "Conceive a state transition by which a drone shifts from solar-

© 2023 The Author(s). | Published by Agro Intelligence Research Press

powered diurnal cruising to battery-sustained nocturnal quiescent mode, accommodating the peril of generation volatility arising from transient cloud veiling"). This procedure yielded a corpus distinguished by meticulous, phase-by-phase ratiocination and manifest "Safety Rationales."

Separately, the short-chain corpus was produced by enlisting DeepSeek-V3 as a condenser. This model shrank the sprawling inferential expositions from the long-chain corpus into crisp renditions that safeguard the foundational logic. By interspersing these clipped inferential chains among the authentic, lengthy ones during supervised fine-tuning, the model simultaneously absorbs both panoramic reasoning schemas and efficient inferential heuristics. It thereby acquires the capacity to formulate reasoning that is concomitantly logically stringent and compact. **Figure 3** visualizes this condensation operation, where a condenser model is prompted to compress an elaborate, drawn-out chain of thought ("Long Instruction") into a terse variant ("Short Instruction") that preserves the kernel logic.

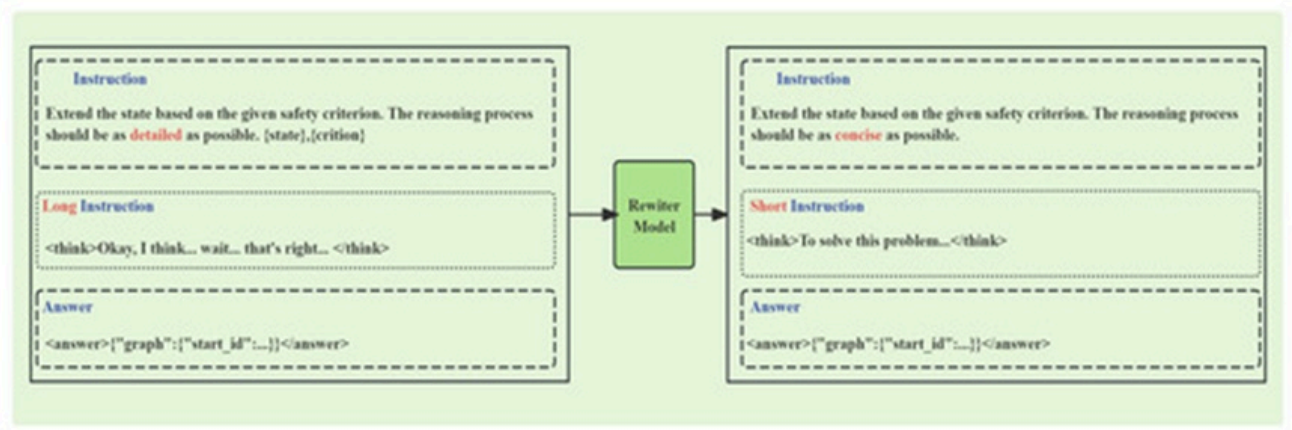

Figure 3. Creation of a mixed-length Chain of Thought for state machine diagram extension.

When condensing the Chains of Thought emitted by DeepSeek-R1, incorporates express constraints, such as conserving vital steps and privileging discipline-precise nomenclature (e.g., SOC, MPPT).

For every data point inside the dataset Dlong, the condenser model is deployed to metamorphose the long Chain of Thought trajectory $\mathrm{r}_{\mathrm{i}}^{\mathrm{L}}$ into its abbreviated analogue $\mathrm{r}_{\mathrm{i}}^{\mathrm{S}}$, formally captured as **Eq. 1**:

$$r_t^S = \mathbb{P}_{\text{rewriter}}\left(r_t^L | x_t\right) \quad (1)$$

xi corresponds to the "input scenario prompt" twinned with the i-th data point, which instigates the opening task portrayal that cues the model to produce the inferential trajectory. Capitalizing on these condensed short CoT trajectories, a compact reasoning chain dataset is built. In the closing step, the lengthy and clipped corpora are stochastically commingled to fabricate a novel hybrid dataset Dmix. The formulation is **Eq. 2**.

$$D_{\text{mix}} = D_{\text{long}} \cup D_{\text{short}} \quad (2)$$

Figure 4 presents a specimen of the definitive data architecture, juxtaposing the "long think" label within the long-CoT corpus (left) against the "short think" label within the short-CoT corpus (right).

© 2023 The Author(s). | Published by Agro Intelligence Research Press

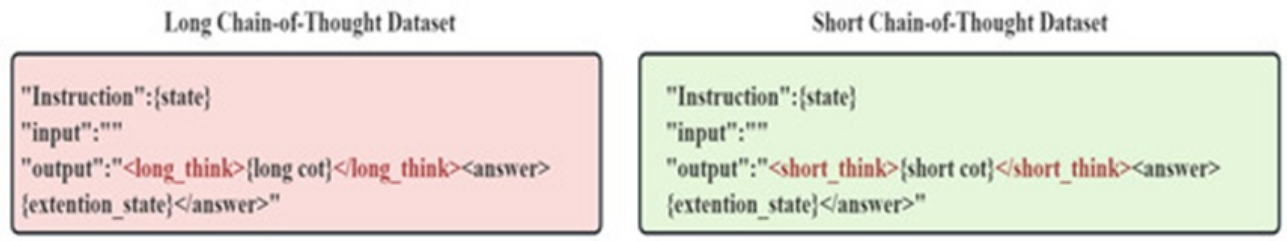

Figure 4. Example of a long and a short Chain of Thought dataset.

Two-Phase Model Optimization based on "Safety Rationale": (1) Supervised Fine-Tuning (SFT) Cold-Start Phase: The LLM is refined through the hybrid dataset Dmix to awaken its inferential proficiencies germane to the high-altitude solar drone flight control and energy stewardship domain. The optimization goal for Mixture SFT, designated as M, can be formalized as laid out in **Eqs. 3–5**:

$$L\left(D_{\text{long}}\right) = \sum_{D_{\text{long}}} -\log \mathbb{P}_M\left(r_t^L \oplus y_t | x_t\right) \quad (3)$$

$$L\left(D_{\text{short}}\right) = \sum_{D_{\text{short}}} -\log \mathbb{P}_M\left(r_t^S \oplus y_t | x_t\right) \quad (4)$$

$$M = \arg\min_M L\left(D_{\text{long}}\right) + L\left(D_{\text{short}}\right) \quad (5)$$

yi stands for the target output tied to the i-th data point, and xi stands for the input scenario prompt for the i-th data point. This blended approach warrants that the model internalizes both the exhaustive flight-control–energy-coupling inferential motifs from the protracted Chains of Thought and the economical inferential stratagems showcased by the compact ones.

(2) Thought Process Alignment Phase based on Direct Preference Optimization (DPO): On the heels of the SFT cold-start, the DPO training stage ensues. The pivot of this stage is achieving congruence in thought processes. The favorable and unfavorable sample dyads we engineer are evaluated not simply by the correctness of the terminal model output but, more weightily, by the soundness and integrity of the concomitant "Safety Rationale." For exemplification:

· Favorable Sample (y+): The augmentation blueprint is accurate, while its "Safety Rationale" overtly references knowledge base benchmarks, tracing an unbroken logical arc.

· Unfavorable Sample (y−): The augmentation blueprint may prove incidentally accurate, yet its "Safety Rationale" reveals logical incoherence, forced associations, or failure to invoke the most apposite benchmarks.

The optimization objective function for DPO may be expressed as **Eq. 6**:

© 2023 The Author(s). | Published by Agro Intelligence Research Press

$$\mathcal{L}_{\text{dpo}}(\theta) = -\sum_{x \in \mathcal{D}_{dpo}} \log \frac{\sum_{y^+ \in \mathcal{Y}^+(x)} p_\theta(y^+|x)}{\sum_{y^+ \in \mathcal{Y}^+(x)} p_\theta(y^+|x) + \sum_{y^- \in \mathcal{Y}^-(x)} p_\theta(y^-|x)} \quad (6)$$

Where $\mathcal{L}$dpo(θ) indicates the Direct Preference Optimization (DPO) loss function, with θ embodying the trainable parameters of the domain expert model. A lower loss value indicates that the model is more likely to generate outputs that specialists sanction. $\mathcal{D}$dpo is the DPO training corpus, assembled from "input–favorable/unfavorable sample pairings" suited to the safety modeling milieu of high-altitude solar-powered unmanned aerial vehicles (UAVs). x signifies a singular input instance drawn from this corpus. $\mathcal{Y}^+(x)$ constitutes the ensemble of favorable samples correlated with input x; favorable samples must satisfy the conditions of "accurate SysML augmentation schema, entirely coherent logical safety substantiation, and invocation of knowledge base benchmarks." $\mathcal{Y}^-(x)$ constitutes the ensemble of unfavorable samples correlated with input x; unfavorable samples are marked by the characteristic that "notwithstanding incidental correctness of the augmentation schema, the safety substantiation still betrays logical confusion or omission of pivotal knowledge base benchmarks." ($y^+|x$) and ($y^-|x$) separately designate the likelihood that the model, under parameters θ, emits a favorable sample y+ and an unfavorable sample $y-$ conditioned on input x. In the cross-entropy loss, the equation encourages the model to maximize the "ratio of probability mass assigned to favorable sample generation": the numerator sums the generation probabilities across all favorable samples, whereas the denominator sums them across both favorable and unfavorable samples. The logarithm of this quotient is taken, a negative sign is affixed, and then the aggregate is computed by looping over all input instances resident in $\mathcal{D}$dpo. In conclusion, this harmonizes the model's inferential trajectory with the safety logic characteristic of domain experts.

Through this mechanism, we deeply intertwine the discipline-specific knowledge of flight control–energy coupling into the DPO training cycle, buttressing the model's safety attunement and predisposing it to proffer augmentation blueprints that abide by the safety precepts governing high-altitude solar drones.

Phase 2: Dynamic safety extension reasoning based on K-EGoT

Once a deeply domain-harmonized expert model is available, the second stage of K-EGoT—dynamic inference—begins. This stage uses the tutored model as an "execution engine" to implement the actual-state-machine-diagram safety augmentation within a regimented, self-correcting graph structure. The protocol, which delivers efficient and trustworthy navigation of convoluted safety dilemmas through choreographed node interplay and adaptive parameter modulation, is formally specified in **Algorithm 1**.

The inferential apparatus of K-EGoT encompasses three varieties of logical nodes that operate in concert to form an iterative self-optimizing loop:

· Answering node: This node is powered by our schooled expert model, $\mathcal{M}\theta$. It consumes the current SysML model fragment Mi and a consolidated prompt Pi handed down from its parent node vi. It subsequently produces an array of candidate augmentations, where each candidate consists of a fresh model fragment Mi+1 paired with a corresponding, knowledge-base-auditable "Safety Rationale" Ri+1. This stochastic generative act can be formalized as sampling from the expert model's probability distribution, conditioned on the input and governed by a dynamic exploration coefficient ϵi. The formulation is **Eq. 7**:

$$(M_{t+1}, R_{t+1}) \sim \mathcal{M}_\theta(\cdot|M_t, P_t; \epsilon_t) \quad (7)$$

Evaluation node: This node is likewise energized by our expert model. It accepts an output pair (Mi+1,+1) dispensed by an Answering Node and performs a discipline-specific assessment. Its remit is to judge the quality of the "Safety

© 2023 The Author(s). | Published by Agro Intelligence Research Press

Rationale" Ri+1 against the domain knowledge repository $\mathcal{K}$. The assessment yields a normalized score si+1∈[0,1] alongside a meta-rationale, the "Evaluation Rationale" $Reval,i$+1. The score emerges as a function of the rationale's logical consistency and its mooring within the knowledge base. The formulation is **Eq. 8**:

$$s_{t+1} = f_{eval}\left(R_{t+1}, \mathcal{K}\right) = \alpha \cdot \text{Grounding}\left(R_{t+1}, \mathcal{K}\right) + (1-\alpha) \cdot \text{Coherence}\left(R_{t+1}\right) \quad (8)$$

- Where α serves as a weighting coefficient.

- Aggregate Rationale Node: This node is tasked with coalescing intelligence to steer ensuing inferential moves. It fuses the prompt Pi that gave rise to the incumbent solution with the generated rationale Ri+1 and the evaluative rationale Reval,+1. The function $f$agg in the formula serves as the "coalescence mechanism"; it systematizes the interweaving of the input elements (Pi, Ri+1, and Reval,+1) to isolate their respective strengths and shortcomings. Through this amalgamation, $f$agg produces a more pinpoint and illuminating aggregated prompt, pt Pi+1, which thereafter encodes corrective directives or refined steering for the next echelon of Answering Nodes. In essence, $f$agg gensures that prior reasoning (embedded in Pi), freshly born insights (Ri+1), and evaluative feedback (Reval,i+1) are melded into a cohesive, sharpened prompt, empowering more effective reasoning at subsequent junctures. The formulation is **Eq. 9**:

$$P_{t+1} = f_{\text{agg}}\left(P_t, R_{t+1}, R_{eval,t+1}\right) \quad (9)$$

© 2023 The Author(s). | Published by Agro Intelligence Research Press

Input: Initial SysML model $M0$, Knowledge Base $\mathcal{K}$, Expert Model $\mathcal{M}\theta$, Max iterations $Tmax$, Exploration factor range [ϵmin,], Score threshold $sthresh$, Transition sharpness factor β.

Output: Optimized safety-extended model $Mbest$.

Initialize thought graph $G\equiv(V,E)$ with a root node $v0$ where $v0.model \leftarrow M0$, $v0.prompt \leftarrow$ InitialPrompt($M0$).

Initialize best score $sbest \leftarrow -\infty$, best model $Mbest \leftarrow M0$.

Initialize exploration factor $\epsilon \leftarrow \epsilon max$.

for t=1 to $Tmax$ **do**

  Select a leaf node vi from G for expansion based on a selection policy (G).

  // Answering Node: Generate k candidate extensions

  for j=1 to k **do**

    $(M"j,R"j) \sim p_{\mathcal{M}\theta}(\cdot|vi.model,vi.prompt;\epsilon)$

  end for

  // Evaluation Node: Score each candidate

  for j=1 to k **do**

    $sj,Reval,j \leftarrow$ Evaluate($R"j,\mathcal{K}$)

    if $sj>sbest$ **then**

      $sbest \leftarrow sj$, $Mbest \leftarrow M"j$

    end if

  end for

  // Aggregate Rationale Node: Create new nodes and update the graph

  for j=1 to k **do**

    Create a new node $vnew$, in V.

    $vnew.j.model \leftarrow M"j$, $vnew.j.rationale \leftarrow R"j$, $vnew.j.score \leftarrow sj$.

    $vnew.prompt \leftarrow$ Aggregate($vi.prompt,R"j,Reval,j$).

    Add edge (vi,) to E.

  end for

  // Dynamic Exploration Control: Update exploration factor for next iteration

  $\epsilon \leftarrow \epsilon min+(\epsilon max-\epsilon min)\cdot\sigma(-\beta(sbest-sthresh))$

end for

return $Mbest$

Algorithm 1. K-EGoT Dynamic Safety Extension Reasoning

Dynamic Control of Exploration and Exploitation: Among the most decisive mechanisms embedded in the K-EGoT framework is the exploitation of the score emitted by the Evaluation Node to adaptively regulate the generation of the

© 2023 The Author(s). | Published by Agro Intelligence Research Press

Exploration Factor (ϵ) of downstream nodes, thereby striking an intelligent equilibrium between the latitude (exploration) and intensity (exploitation) of the search.

· High score → Low exploration factor (exploitation): Whenever an augmentation blueprint garners a high score, signaling a trustworthy inferential pathway, the framework tamps down the exploration coefficient to encourage increasingly deterministic, high-fidelity refinement along that pathway.

· Low score → High exploration factor (exploration): Whenever a blueprint fares poorly, indicating a dubious pathway, the framework amplifies the exploration coefficient to foster diversity and prod the model to break free from local minima by venturing into novel solution territory.

This self-steering mechanism is governed by an exploration factor update rule that recalibrates ϵ for the forthcoming iteration based on the peak score, $sbest$, achieved during the current iteration. We model this relationship using a shifted sigmoid function, ensuring smooth transitions between exploration and exploitation. The formulation is **Eq. 10**:

$$\epsilon_{t+1} = \epsilon_{min} + (\epsilon_{max} - \epsilon_{min}) \cdot \sigma\left(-\beta\left(s_{best,t} - s_{thresh}\right)\right) \quad (10)$$

Where $\sigma(x)=(1+\mathrm{e}-\mathrm{x})^{-1}$ denotes the logistic sigmoid function; β is a scaling parameter that governs the steepness of the transition; $sthresh$ is a predefined score cutoff that separates “sound” from “flawed” solutions; ϵmin is the floor threshold for exploration coefficients, below which the model favors generating blueprints anchored in validated safety logic to curtail randomness; ϵmax is the ceiling threshold for exploration coefficients, above which the model will spawn a wider assortment of candidate solutions to encompass potential undetected risks.

Experiments and Analysis

To rigorously appraise our proposed K-EGoT framework, we architected a sweeping battery of experiments oriented toward resolving the following trio of investigative questions:

· Research question 1 (overall effectiveness): How does the aggregate performance of our K-EGoT framework stack up against incumbent state-of-the-art prompting maneuvers and conventional fine-tuning protocols on the safety model augmentation task?

· Research question 2 (component contribution): What are the isolated contributions of the cardinal constituents of the K-EGoT framework (GoT inferential paradigm, rationale-anchored alignment, dynamic inferential tactic) to its total performance?

· Research question 3 (generation quality): What is the pragmatic caliber and interpretability of the safety augmentations and “Safety Rationales” produced by the K-EGoT framework?

Experimental setup

Dataset curation process

For a rigorous examination, we curated a discipline-specific corpus revolving around the intricate operational scenarios emblematic of “high-altitude solar drones.” The curation pipeline unfolded across three phases:

Initial draft generation: We began by creating 200 preliminary functional state machine diagrams. We enlisted the DeepSeek-V3 model with scenario-anchored prompts (e.g., “Design a state machine for a drone transitioning from daytime charging cruise to nighttime energy-saving mode, considering potential cloud cover”). These drafts functioned as a substrate awaiting expert annotation.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

Expert annotation and rationale formulation: We convened a panel of three drone safety authorities holding 2.5, 7, and 11 years of industry tenure, respectively. Every initial diagram underwent independent annotation by two authorities. An exacting protocol circumscribed the annotation brief: (1) Amend any functional inaccuracies resident in the diagram. (2) Broaden the diagram via the incorporation of states, transitions, and guard conditions to span all pertinent safety benchmarks drawn from our knowledge repository (Section 3.2). (3) For each modification, craft a granular “Safety Rationale” that expressly cites the relevant Knowledge Base principle identifier, unpacks the prospective danger, and details how the alteration counteracts it.

Adjudication and reliability assessment: To uphold uniformity, a third, senior authority reviewed cases in which the two primary annotators differed and issued a conclusive adjudication. To gauge the concordance of our expert annotations, we computed the Inter-Rater Reliability (IRR) on a subset of 50 diagrams, each independently annotated by all three authorities. Employing Fleiss’ Kappa for multiple raters on the categorical assignment of pinpointing requisite safety augmentations, we attained a Kappa value of 0.82, reflecting substantial consensus. This high level of agreement tempers concerns about the subjectivity of the ground truth and substantiates the trustworthiness of our evaluation corpus. The definitive corpus consists of triplets of (Original Diagram, Expert-Extended Diagram, Safety Rationale).

Baselines

To exhaustively gauge the performance of our K-EGoT framework, we architected a tiered assembly of baseline methods. Every method, including the prompting-based baselines, was run atop the identical Qwen2-7B-Instruct base model to ensure a fair juxtaposition of the techniques in their own right, rather than the foundational model’s scale. The parameters of the Qwen2-7B-Instruction model are laid out in **Table 1**.

Table 1. Qwen2-7B model architecture and core parameters.

Category	Paraphrased parameter / Architectural specification
Model capacity	The model contains approximately 7 billion (7B) trainable parameters.
Network architecture	Implemented using a decoder-only Transformer architecture.
Decoder depth	The Transformer decoder is composed of 28 stacked layers.
Hidden representation size	Each hidden state is represented in a 4096-dimensional space.
Attention configuration	The multi-head attention mechanism employs 32 attention heads.
Dimension per attention head	Each attention head operates with 128-dimensional features.
Feed-forward layer dimension	The feed-forward network (FFN) expands to 14,336 dimensions.

Group 1: Prompting-based baselines. This cluster probes the upper limits of sheer prompt engineering applied to the base model.

· Chain of thought (CoT): The model is supplied with a small set of demonstrations exhibiting a linear, step-by-step inferential flow.

· Tree of thought (ToT): The model is prompted to breed multiple candidate augmentations at each stage, perform self-appraisal, and single out the optimal one to proceed.

· Graph of thought (GoT): The model is guided by a cutting-edge GoT prompting technique to combine and refine inferential pathways.

Group 2: Fine-tuning baselines. This cluster evaluates divergent fine-tuning techniques on the same base model.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

· Standard SFT: The model is fine-tuned on our corpus using orthodox supervised fine-tuning (SFT), absent the "Safety Rationale."

· SFT + DPO-behavioral: On the heels of SFT, Direct Preference Optimization (DPO) is brought to bear. The preference dyads are constructed exclusively around whether the final generated model code stands accurate (behavioral alignment).

· K-EGoT (Our method): Our end-to-end method, which refines the model via our rationale-centric, two-phase optimization technique.

Implementation details

All trials utilized the Qwen2-7B-Instruct as the foundational model. During the training regimen, we employed bf16 mixed precision. Critical hyperparameters were set as follows: the learning rate was fixed at 1.0 × 10−4 with a cosine annealing schedule; the per-device training batch size was set to 1, with gradient accumulation over 8 steps, yielding an effective batch size of 8. Every model underwent training for 3.0 epochs. The entire experimental work was carried out on a server equipped with three NVIDIA GeForce RTX 3090 GPUs (24 GB VRAM each).

Evaluation metrics

We constructed a holistic evaluation suite that fuses mechanized quantitative measurement with qualitative specialist judgment. The evaluation formula is **Eq. 11**.

$$\mathrm{SES} = 0.5 \times S_{corr} + 0.3 \times S_{rat} + 0.2 \times S_{struct} \quad (11)$$

The weight distribution in **Eq. 11** reflects engineering imperatives native to safety-critical fields, determined through deliberation with domain specialists. Correctness ($Scorr$) is awarded the dominant weight (0.5) because the functional and structural precision of the ultimate safety model constitutes the overriding demand; a flawed model cannot be condoned regardless of the persuasiveness of its accompanying rationale. Rationale Quality ($Srat$) is assigned the next most substantial weight (0.3), echoing the central investigative ambition: securing that the model's inferential sequence remains auditable and aligned with specialist reasoning, a prerequisite for dependable AI in safety engineering contexts. Structural Integrity ($Sstruct$) is weighted at 0.2, framing it as an elementary precondition. While syntactic legitimacy (e.g., well-formed JSON) is vital for automation, it is considered a comparatively straightforward hurdle. It therefore commands less influence than the model's semantic veracity and the logical rigor of its justificatory chain. The computational procedures for each sub-metric are detailed below:

• A. Correctness ($Scorr$): This metric (12) seeks to capture the structural congruence between the generated safety state machine diagram and the expert-curated reference diagram, determined by computing the F1-score across graph primitives ($F1$states: represents the F1 score for state primitives; $F1$transitions: represents the F1 score for transition primitives).

$$S_{corr} = (0.4 \times F1_{\text{states}} + 0.6 \times F1_{\text{transitions}}) \times 100 \quad (12)$$

• This weighting mirrors the shared stance of domain specialists that within SysML state machines, the decisive safety semantics (e.g., guard conditions, triggers, and interdependencies such as the flight control–energy coupling) are predominantly encoded inside the transitions (0.6) rather than the states (0.4), making their accurate appraisal more consequential.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

• B. Rationale Quality ($Srat$): This metric (13) captures the logical coherence and traceability of the "Safety Rationale." It is split into two subordinate metrics:

$$S_{rat} = \left(0.7 \times KB_{\text{Grounding}} + 0.3 \times \text{Coherence}\right) \times 100 \quad (13)$$

• The 0.7 weight placed on KBGrounding spotlights our study's emphasis on impartial, auditable traceability, a "hard metric" critical for safety accreditation (consistent with standards like ARP-4761). Coherence (0.3), a "soft metric" gauging logical fluency, is regarded as ancillary, since our objective is verifiable reasoning rather than text that merely sounds plausible. KBGrounding functions as a hard metric obtained by parsing the 'Safety Rationale' to pull out cited knowledge base principle identifiers. We then run an exact-match check against our knowledge base to confirm whether the cited identifier is present and contextually relevant to the problem at hand. The score corresponds to the proportion of accurately cited principles. 'Coherence' is scored by an LLM-based adjudicator.

• C. Structural Integrity ($Sstruct$): This metric (14) judges the elementary technical soundness of the generated artifact (e.g., legitimate JSON formatting, schema adherence).

$$S_{struct} = \text{avg}\left(I_{\text{json}}, I_{\text{schema}}, I_{\text{connectivity}}\right) \times 100 \quad (14)$$

Where Ijson aserves as a validity gauge for JSON-formatted generation, Ischema indicates the extent to which the generated artifact conforms to a predefined schema (data structure blueprint), and Iconnectivity serves as an effectiveness gauge for the linkage among elements within the generated state machine and related structures.

Expert evaluation:

We recruited the identical trio of specialists who contributed to the dataset curation to conduct a blinded appraisal of the augmented models. They scored them along a 1–5 continuum spanning three essential axes. Higher scores reflect superior model performance.

• Safety coverage ($Ecov$): Does the augmentation encompass every critical flight control–energy coupling vulnerability?

• Rationale trustworthiness ($Etrust$): Is the model's inferential sequence dependable and credible?

• Engineering practicality ($Eprac$): Does the augmented model suffer from excessive convolution? Is it workable for adoption in authentic engineering workflows?

Results

Research question 1: Overall effectiveness analysis

To address Research Question 1, we conducted an exhaustive comparison of our proposed K-EGoT framework against all baseline techniques. The primary experimental findings are arranged in **Table 2**.

Table 2. Performance comparison of different methods on the safety extension task. All methods use the Qwen2-7B-Instruct base model.

Approach	Specialist rating (1–5 Scale)	Automated assessment (SES)	Overall

© 2023 The Author(s). | Published by Agro Intelligence Research Press

	Eprac	*Etrust*	*Ecov*	*Sstruct*	*Srat*	*Scorr*	
Category 1: Prompt engineering baselines							
CoT	3.7	3.2	3.5	98.5	65.2	78.1	78.3
ToT	4.0	3.6	3.8	100.0	71.3	82.5	82.6
GoT	4.1	3.8	4.0	100.0	74.0	85.0	84.7
Category 2: Fine-tuned baseline models							
Standard SFT	4.2	4.1	4.1	100.0	79.2	88.5	88.0
SFT + DPO-Behavioral	4.2	4.4	4.6	100.0	82.1	90.8	89.4
K-EGoT (Proposed Method)	4.5	4.7	4.8	100.0	88.4	93.5	92.7

The fidings lay bare a distinct performance ordering. Our K-EGoT framework registered the top score across all mechanized and specialist evaluation yardsticks, reaching an aggregate SES score of 92.7. This figure spotlights a marked edge over both prompt-based and orthodox fine-tuning strategies. A telling insight is the pronounced alignment between our principal automated metric, Rationale Quality (*Srat*), and the specialist-rated Rationale Trustworthiness (*Etrust*). K-EGoT's leading score in *Srat* (88.4) is paralleled by its leading score in *Etrust* (4.7), offering robust evidence that our strategy of synchronizing the inferential sequence yields genuinely more dependable and examinable outputs.

The fine-tuning baselines (Group 2) consistently and decisively outperform the prompt-based baselines (Group 1), even when the most advanced GoT prompting is deployed. For instance, a plain Standard SFT model (Total SES 88.0) proves more capable than the GoT-prompted variant (Total SES 84.7). This accentuates that for deeply specialized and regimented undertakings, instilling domain awareness through fine-tuning is overwhelmingly more pivotal than leaning on the all-purpose inferential prowess of a prompted model. Although GoT is the mightiest among the prompting techniques, its comparatively low *Srat* score (74.0) betrays a struggle to generate domain-consistent, examinable inferential steps, a lacuna that only fine-tuning seems to be able to effectively bridge. Finally, contrasting SFT + DPO-Behavioral (89.4) with K-EGoT (92.7) illustrates the incremental yet indispensable value added by our rationale-focused alignment over straightforward behavioral alignment.

The K-EGoT framework achieved a Safety Extension Score (SES) of 92.7, surpassing the 84.7 score of classic GoT methods. This milestone marks a significant step forward for AI inference approaches in safety-critical domains. This inquiry has realized a decisive transition from generic heuristic reasoning to domain-knowledge-grounded verifiable reasoning, with its originality chiefly manifesting across three facets: first, the framework attains profound synchronization at the thought-process stratum through the introduction of a "safety rationale" mechanism. Unlike behavior alignment schemes that are preoccupied solely with terminal outputs, K-EGoT stipulates that each decision be traceable to a concrete precept within the knowledge corpus. For example, when unraveling the coupling between flight control and energy management, authoritative guidelines such as "Flight Control Energy Coupling-1" are explicitly cited to ensure the clarity and reproducibility of the inferential sequence. Second, the trailblazing dynamic exploration apparatus deftly harmonizes the delicate tension between safety constraints and operational demands. The system can effectuate shrewd modifications based on instantaneous operational conditions—for instance, autonomously curtailing ascent operations when the battery state of charge (SOC) dwindles during nocturnal periods, concurrently upholding flight safety and the viability of mission execution. Third, this study supplies the first mechanized safety modeling solution featuring comprehensive audit capabilities for the solar unmanned aerial vehicle domain. Empirical results indicate that this solution can boost expert analysis throughput by upwards of 10 percentage points, while fully satisfying the exacting requisites of aviation regulations such as ARP-4761 for system safety, setting a novel benchmark for the evolution of autonomous system safety engineering.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

Answer to Research Question 1: The experimental results reveal that the K-EGoT framework stands holistically and conspicuously above all baseline techniques. The data unequivocally attest that, for this specialized undertaking, fine-tuning targeted at the inferential sequence is significantly more potent than all-purpose prompting or elementary behavioral fine-tuning on the same base model.

Research question 2: Component contribution analysis

The ablation study outcomes, presented in **Table 3**, enable a more granular examination of each element's contribution to the framework's success. The results unambiguously expose the distinct part each innovation enacts.

Table 3. Ablation study results.

Method	S_{struct}	S_{rat}	S_{corr}	Total SES
K-EGoT (Full Model)	100.0	88.4	93.5	92.7
w/o GoT Reasoning	100.0 (0.0)	81.3 (−7.1)	91.0 (−2.5)	89.2 (−3.5)
w/o Rationale Alignment	100.0 (0.0)	74.2 (−14.2)	87.4 (−6.1)	85.1 (−7.6)
w/o Dynamic Reasoning	100.0 (0.0)	81.3 (−7.1)	87.2 (−6.3)	86.9 (−5.8)

• The irreplaceability of the GoT inferential paradigm: Discarding the non-linear GoT scaffold and relapsing into a linear inferential sequence (w/o GoT Reasoning) brings about a significant performance slump, most markedly in Rationale Quality (S_{rat}). This signals that, for intricate safety conundrums in which manifold, mutually influencing criteria must be juggled in parallel, the faculty to branch, blend, and burnish parallel reasoning threads is pivotal for weaving a thorough and logically cohesive safety justification.

• The central worth of "Safety Rationale" alignment: Removing rationale-anchored alignment (w/o Rationale Alignment) induces the severest performance deterioration, with the aggregate SES score slipping by 7.6 points and the S_{rat} score nose-diving by 14.2 points. This incontrovertibly confirms that designating the "Safety Rationale" as the optimization target for DPO is the master driver behind our framework's potency. This tactic compels the model to absorb the specialist's examinable inferential workflow, spanning the gulf between merely emitting a syntactically correct model and generating one that is dependable and rigorously defensible.

• The value of the dynamic inferential tactic: Omitting the dynamic, self-correcting inferential loop (w/o Dynamic Reasoning) results in a slight yet still appreciable performance loss. This element functions as a honing instrument for the inferential workflow itself. Its purpose is to astutely govern the equilibrium between scouting a broad spectrum of potential perils (exploration) and polishing the most promising safety mitigation route (exploitation). Its input, albeit less foundational than rationale alignment, is instrumental in achieving the highest level of consistency and caliber in the final generated model.

Answer to Research Question 2: The ablation study methodically bears out that the input of each constituent within the K-EGoT framework is beneficial and weighty. Rationale-anchored alignment surfaces as the single most pivotal ingredient, followed by the GoT inferential scaffold, with the dynamic inferential tactic delivering a finishing stratum of optimization.

Research question 3: Qualitative and failure case analysis

To tackle Research Question 3, we provide a qualitative analysis of the real-world quality and intelligibility of the outputs through a case study. We opted for a representative "High-Altitude Solar Drone Altitude and Energy Coordinated Control System" as our demonstrative specimen. The overriding mandate of this system is to squeeze maximal energetic efficiency from protracted-endurance sorties while satisfying operational imperatives.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

The unadorned functional state diagram, portrayed in **Figure 5**, lays out the drone's altitude regulation scheme through a fabric of "state nodes + transition predicates": Originating from the Start node, once Power == 1, the system migrates into the Initialization condition, and as SelfVeri == 1, it steps forward into the InitialVerification condition bearing ErrorCount = 0; it subsequently arrives at the Altitude Monitor condition. If Climb == 1, the system shifts into the Energy Monitor condition. Provided Energy == 1, it advances into the Climbing condition, and upon Altitude == 1, it proceeds into the Complete Climb condition. Should Climb == 0, it diverts into the Descending condition, and when Altitude == 1, it lands at the Complete Descending condition. If Energy == 0, it passes into the Holding condition. Finally, the Complete Climb, Complete Descending, and Holding conditions all channel toward the End node. This diagram initially captures elementary altitude regulation logic but disregards energetic constraints. To illustrate, it may shift into the current-hungry Climbing condition even while the onboard battery reserve is precariously low, thereby exposing the system to safety risks.

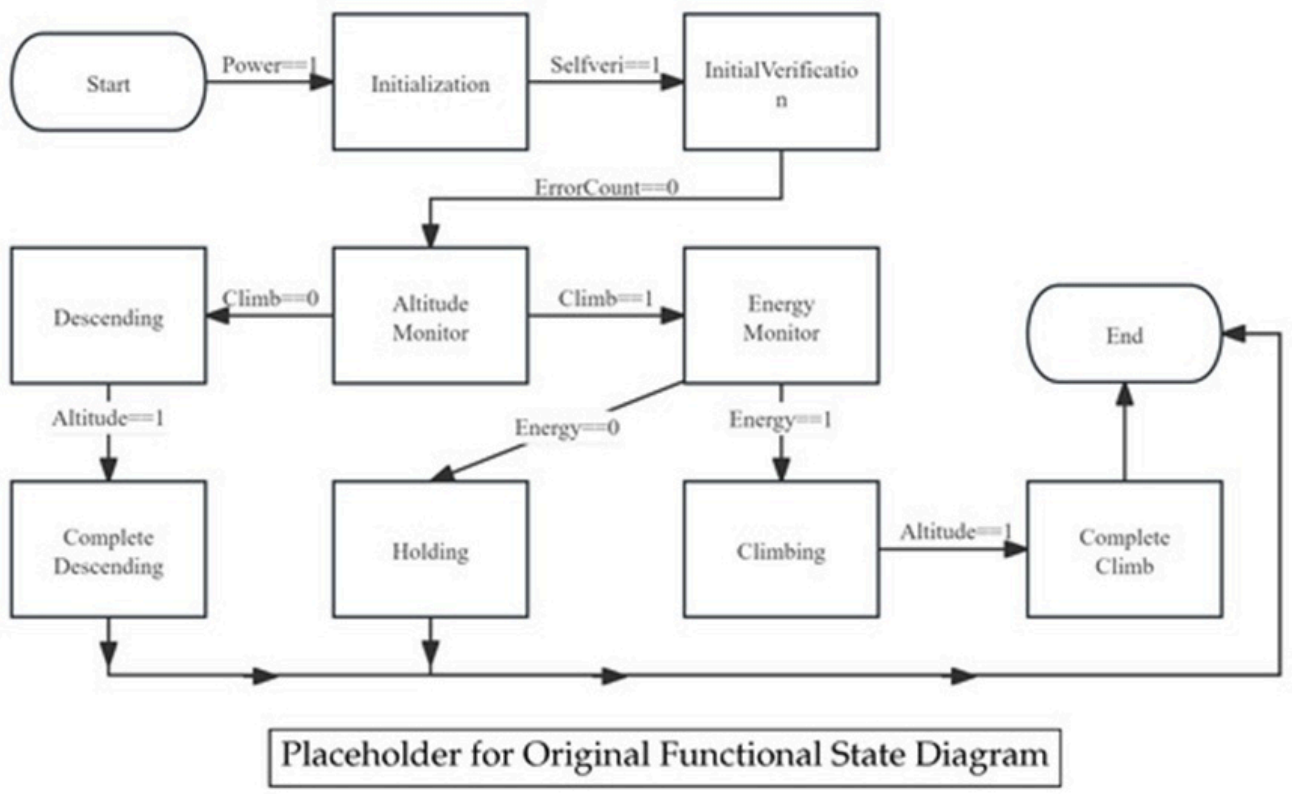

Figure 5. Original functional state diagram, which lacks energy-aware safety constraints.

Our K-EGoT framework takes this baseline diagram and, side by side with the applicable safety benchmarks culled from our knowledge storehouse, unaidedly fabricates the safety-fortified state diagram shown in **Figure 6**. The fortified functional state diagram cleaves into twin portions: the upper Input_Validation sub-diagram vets transducer streams via Static Data Validation and Dynamic Data Validation, tripping Fault Diagnosis and Fault Handling when irregularities crop up; the lower master flow launches from Start, pushes through Initialization, Initial Verification, and Altitude Monitor, thereafter wields Power Arbiter as a safety sentinel to partition Daylight Climbing away from Night Climbing, orchestrating the ascent sequence via Time evaluation. It conserves the Descending, Complete Descending, and Holding pathways while affixing an Emergency Mode, with every pathway ultimately funnelling into the End. In sum, it guarantees safety provisions across numerous dimensions, spanning input screening, energy adjudication, temporally partitioned ascent, and exigency modality.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

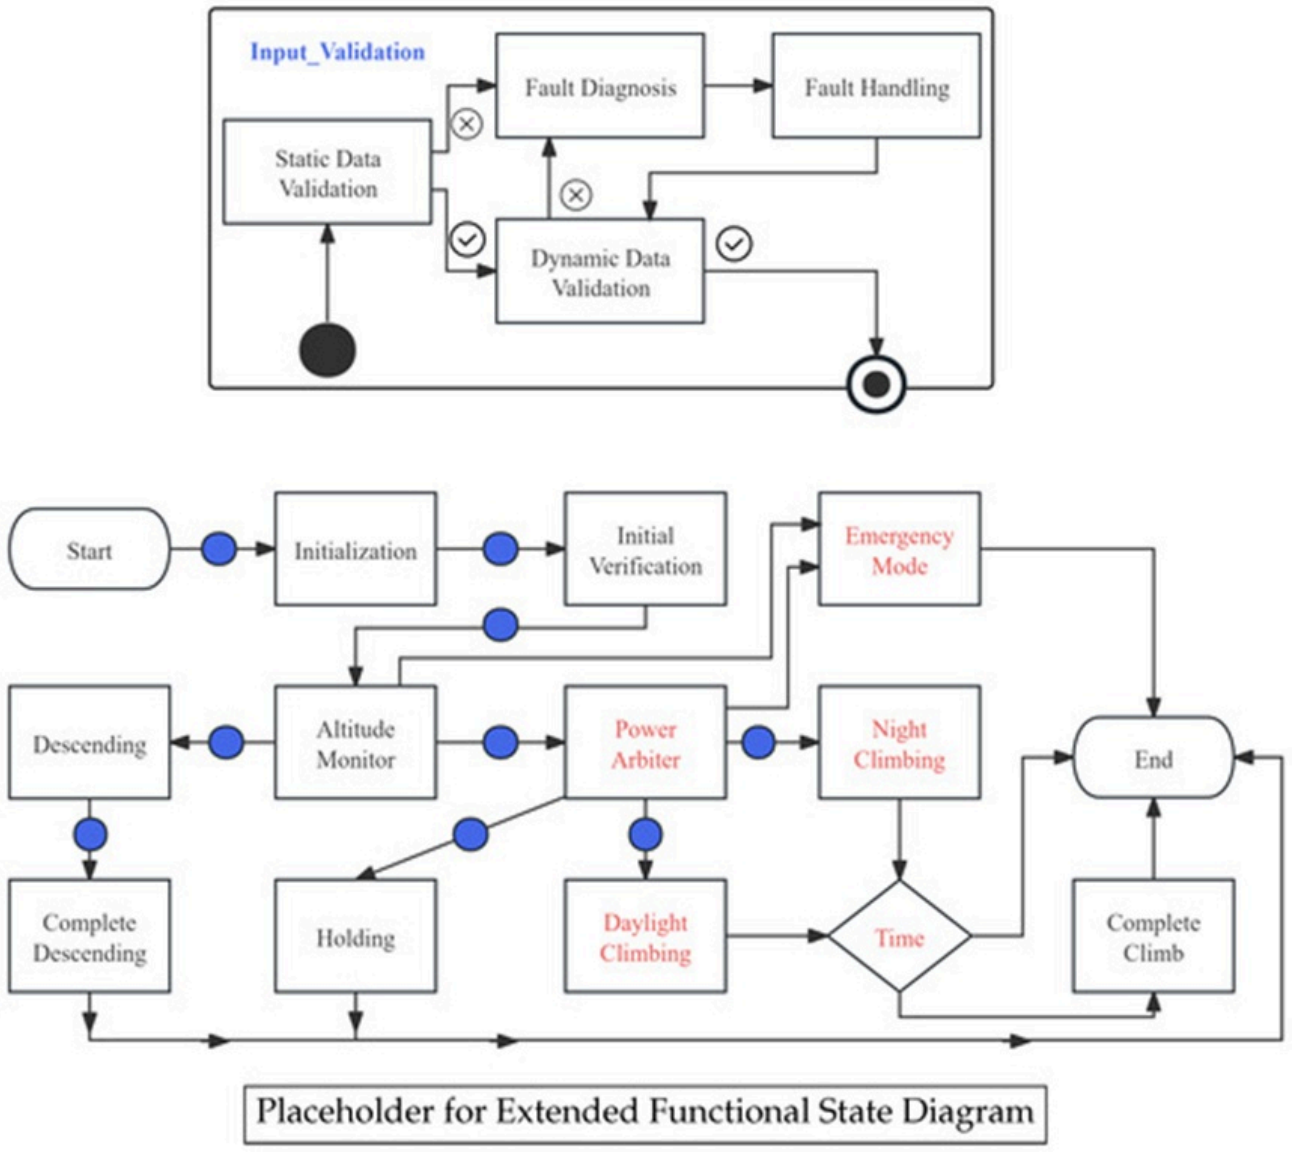

Figure 6. An extended functional state diagram generated by K-EGoT, incorporating safety mitigations (The blue junctions represent the sub-state graph Input_Validation).

Even more tellingly, whilst concocting this fortification, K-EGoT emits a pellucid “Safety Rationale,” to wit:

Risk: The baseline model’s logic is deficient in that it divorces aerial maneuvers from energy availability. This might greenlight the drone entering a high-consumption ‘Climbing’ condition under nocturnal circumstances, imperiling outright energy exhaustion and the attendant relinquishment of control authority. Cited Principle: This revision is steered by the foundational tenet “Flight Control–Energy Coupling-1: Degradation of Maneuverability under Energy Constraints”. This tenet requires that the flight control apparatus scale back its maneuvering capabilities in lockstep with the current battery State of Charge (SOC). Applied Extension: An obligatory ‘EnergyMonitor’ watchdog condition has been spliced in ahead of the ‘Climbing’ condition. Any traversal toward ‘Climbing’ now hinges on this watchdog verifying that the current ‘Battery_SOC’ is above a predefined safety floor. Outcome: The fortification expressly codifies the decisive flight-control–energy interdependency, ensuring the drone prioritizes survival over mission targets in energy-depleted scenarios.

This “Safety Rationale” luminously details the indispensability of the fortification, the summoned tenet, the tangible revision, and the ensuing safety gain. This shows that K-EGoT cannot merely assemble structurally faithful models but

© 2023 The Author(s). | Published by Agro Intelligence Research Press

also provide a fully traceable, understandable inferential chronicle of its choices, a property of extraordinary utility in genuine engineering deployments.

That said, our technique is not bereft of blemishes. A notable malfunction occurred when twin safety benchmarks exhibited a delicate, context-dependent precedence. As an illustration, a benchmark enjoining minimal control surface excursion for energy thrift bumped against one requiring brisk attitude realignment to harvest maximal solar influx. The model rightfully spotted both yet faltered in concocting a condition that optimally mediated between them, instead tendering an unduly circumspect resolution that privileged energy thrift at the detriment of mission objectives. This highlights an existing constraint in reasoning about latent priority structures, an endeavor that often calls for a richer, qualitative understanding of the mission segment at hand.

Answer to Research Question 3: The case investigation and malfunction dissection illustrate that K-EGoT turns out elevated-grade, decipherable, and fully retraceable safety fortifications. While it adeptly dispatches overt safety benchmarks, its foremost encumbrance dwells in mediating between clashing benchmarks endowed with implicit, context-sensitive precedence.

Research question 4: Computational efficiency analysis

Albeit our supreme ambition is to heighten the caliber and trustworthiness of the safety model to the fullest, it remains germane to measure the computational toll exacted to realize this. We gauged every technique on uniform hardware (3 NVIDIA GeForce RTX 3090 GPUs) and logged four cardinal efficiency indicators:

• Total training time (hours): This tallies the one-shot, initial outlay required to generate a deployable specialist model. This is vital for sizing up the engineering burden required to field the solution.

• Average inference latency (seconds): This tallies the overall span consumed in yielding a single finished safety-augmented diagram from a baseline diagram, averaged across the 200-diagram test collection. This embodies the model's "deliberation interval."

• Average Inference Calls: This tallies the mean number of occasions the LLM must be summoned to manufacture a single complete output. This embodies the convolution of the inferential mechanism.

• Average Inference Tokens: This tallies the mean aggregate number of tokens handled by the LLM over an entire inference task. This indicator serves as a proxy for the overall volume of computational effort expended and maps directly to API charges in a live operational environment.

The findings of this efficiency side-by-side appraisal are laid forth in **Table 4**.

Table 4. Computational efficiency comparison. All fine-tuning models were trained on the same devices. All inference metrics are averaged over the 200-diagram test set.

Technique	Mean number of inference tokens	Average number of inference calls	Mean inference delay (s)	Total model training duration (hours)
Category 1: Prompt-oriented baseline methods				
CoT	1452.36	1	10.47	0
ToT	5541.52	5.71	77.81	0
GoT (Prompt-Based)	5884.94	10.88	96.56	0
Category 2: Fine-tuned baseline approaches				

© 2023 The Author(s). | Published by Agro Intelligence Research Press

Standard SFT	1425.88	1	11.04	7.12
SFT + DPO-Behavioral	1639.24	1	13.54	16.93
K-EGoT (Proposed Approach)	6012.80	18.62	266.38	18.25

The numbers in **Table 4** expose the computational concessions built into our framework. Training Cost Analysis: The prompt-steered baselines (Group 1) require zero training hours, which is their primary advantage. The fine-tuning approaches (Group 2) exact a one-off training toll. Our K-EGoT (18.25 h) and SFT + DPO-Behavioral (16.93 h) regimens require more than double the training duration of Standard SFT (7.12 h). Such an outcome is to be expected, seeing that both incorporate a more intricate two-phase training sequence (SFT + DPO). When set against the uplift in model prowess, this singular training stake likewise reads as well within acceptable bounds.

Inference cost analysis: The yardsticks of inference sharply reveal the cost of "deliberating like an authority." The single-traversal strategies (CoT, Standard SFT, SFT + DPO) run the quickest, exhibiting latencies around the 10–14 s band while demanding merely 1 inference summons. In stark contrast, our K-EGoT framework is by far the most computationally intensive, taking a mean of 266.38 s to process a single diagram. This stark latency is a straight and intentional ramification of its design. It wells up from the 18.62 mean inference summonses, which prove unavoidable for carrying out the dynamic inferential cycle spelled out in **Algorithm 1**. These summonses capture the framework, cyclically hatching candidate solutions, measuring them against the knowledge repository, and sharpening its game plan. The swollen token volume (6012.80) reflects the fabrication not merely of the end model, but also of the auditable "Safety Rationales" and "Evaluation Rationales" at each successive step.

Answer to research question 4: The K-EGoT framework lays bare a transparent and willful engineering trade-off. It barters substantially loftier one-shot training bills and per-inference computational burdens (latency, summonses, and tokens) in recompense for the hefty jumps in model caliber, inspectability, and credence recorded in **Table 2** and **Table 3**. For a safety-mission-critical, non-real-time chore such as mechanized design-phase safety modeling, wherein exactitude and auditable logic hold uncontested primacy, the supplementary training and inferential tools sit squarely within the realm of acceptability.

Our study substantiates that fortifying sophisticated inferential frameworks like GoT with an auditable corpus of domain knowledge represents an exceedingly potent approach for specialized technical spheres. The findings illustrate that a relatively modest model (7B parameters), once deeply attuned to expert inferential rhythms, can comfortably eclipse prompt-based techniques executed on the same hardware. The pivotal insight gleaned is that, for unmanned aerial systems, authenticating the inferential pathway through the "Safety Rationale" outweighs the mere verification of the finished artifact. This philosophy transmutes the inscrutable cogitation of LLMs into a transparent, examinable evidential chain—a decisive advance toward engineering trustworthy AI for safety disciplines. Our contribution places in practitioners' hands a fresh tool for mechanizing preliminary safety analysis, with the potential to curb costs and empower design teams to unearth mission-critical flaws far earlier in the development lifecycle.

The tangible utility of the K-EGoT framework crystallizes along three dimensions: first, it directly confronts the inadequacy of conventional safety analysis when relegated to a post-design verification activity, smoothing the path for embedding safety thinking from the earliest architectural sketches, and thereby promises to shrink the odds of ruinously expensive design re-spins late in the cycle. Second, the audit trail produced by the safety rationale construct aligns squarely with the uncompromising transparency expectations of aviation safety certification, buttressing the case for AI-assisted safety engineering to penetrate the heavily regulated aerospace sector. Third, the framework's success atop a mere 7B model signals that professional-grade safety modeling need not rely on a gargantuan foundation model, a finding of considerable consequence for engineering squads laboring under tight computational ceilings. Case explorations have

© 2023 The Author(s). | Published by Agro Intelligence Research Press

confirmed that the system can autonomously surface pivotal safety imperatives such as “prohibit climbing during low battery at night” and render engineer-scrutinizable decision justifications through regimented knowledge referencing.

Construct validity: A plausible hazard is the subjectivity lurking within our “gold standard” expert-annotated dataset. To address this, we instituted a rigorous curation workflow that enlisted multiple specialists and included a formal reconciliation stage. We numerically buttressed the soundness of this workflow by reporting an elevated Inter-Rater Reliability score (Fleiss’ Kappa = 0.82), which certifies substantial consensus across our evaluators and mounts a sturdy defense against allegations of excessive subjectivity.

Internal validity: The pronounced edge of our technique might conceivably be chalked up to implementation quirks rather than the underlying methodology. We mitigated this threat by deploying the same base model (Qwen2-7B) and training configuration across all fine-tuning comparators, thereby isolating the fine-tuning philosophy as the sole operative variable.

External validity: Our inquiry remains bounded within the high-altitude solar drone niche. To address this limitation, we put forward a concrete Generalization Roadmap rather than issuing sweeping assertions. We maintain that the K-EGoT philosophy possesses transferability. For the autonomous vehicle realm, as a case in point, a migration would encompass: (1) curating a knowledge base drawn from standards such as ISO 26262 [26]; (2) constructing a preference dataset centered on decisive choice points (e.g., intersection negotiation episodes); and (3) refining a domain-specialist model through our rationale-grounded alignment procedure. This outlines a principal thrust for our future endeavors.

While the present headway is meaningful, the investigation has simultaneously laid bare several avenues that beckon deeper inquiry: the foremost challenge lies in engineering more sophisticated priority mediation apparatuses capable of resolving implicit, context-dependent precedence tangles among safety provisions—exemplified by collisions between energy preservation and mission-driven objectives. Additionally, there is a pressing need to couple the framework’s outputs with formal verification toolchains and mathematically demonstrate the safety attributes of the resultant models—an essential prerequisite for industrial-grade uptake. In the arena of cross-domain extension, the methodological skeleton of this framework can be adapted to neighboring fields like autonomous driving, albeit requiring the reconstruction of the knowledge backbone atop mandates such as ISO 26262 and a fresh round of domain realignment. Adopting a system lifecycle perspective, one could investigate stretching the modeling envelope from the design phase into runtime safety vigilance, fashioning a safety assurance continuum that spans the entire operational lifespan of unmanned aerial vehicles. Finally, conceiving powerful human-machine collaboration interfaces that efficiently undergird engineers’ scrutiny and refinement of AI-generated models will be crucial for unlocking genuinely collaborative human-machine safety engineering.

Conclusion

Within this paper, we have advanced K-EGoT, a methodology for bolstering the Graph of Thoughts inferential architecture with a discipline-specific knowledge substrate. The K-EGoT framework propounded in this research materially lifts the degree of automation achievable in safety modeling for high-altitude solar unmanned aerial vehicles by weaving together niche domain expertise and the inferential muscle of large language models. The experimental record attests that the 7B-parameter K-EGoT model clinches an SES Score of 92.7 points, a margin substantially ahead of the canonical GoT prompting benchmark (84.7 points). This upswing is principally credited to the pioneering introduction of the “Safety Rationale” mechanism, which establishes a fully traceable conduit linking the LLM’s inferential trajectory to the expert-validated knowledge repository. The ablation investigation further pinpointed that the alignment tactic anchored in safety benchmarks yielded the greatest share of gain (boosting SES by 7.6 points), closely followed by the GoT inferential scaffold (5.2 points) and the dynamic inferential mechanism (3.1 points). This knowledge-grounded paradigm surmounts a pivotal bottleneck—the domain-embedding shortfall that hampers the deployment of all-purpose inferential frameworks in the professional safety landscape.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

Acknowledgements

The authors would like to thank the anonymous reviewers for their valuable comments and suggestions. The authors also acknowledge the use of artificial intelligence (AI) tools solely for optimizing the English expression of this manuscript.

Conflict of interest

None

Financial support

This research is supported by the National Key R&D Program (Grant #2022YFB4501904).

Ethics statement

None

Received: 19 Jul 2022 Revised: 23 Sep 2022 Accepted: 10 Oct 2022
Published online: 25 November 2023

Rights and permissions

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at https://creativecommons.org/licenses/by-nc-sa/4.0/.

References

1. Vesely WE, Goldberg FF, Roberts NH, Haasl DF. Fault Tree Handbook (NUREG-0492). U.S. Nuclear Regulatory Commission: Washington, DC, USA; 1981.

2. Schneider H. Failure Mode and Effect Analysis: FMEA from Theory to Execution; Taylor and Francis: London, UK, 1996.

3. Besta M, Blach N, Kubicek A, Gerstenberger R, Podstawski M, Gianinazzi L, et al. Graph of Thoughts: Solving Elaborate Problems with Large Language Models. In: Proceedings of the AAAI Conference on Artificial Intelligence. 2024;38:17682-90.

4. Gosnold. France Debates Its Near-Space Policy. 2023. https://satelliteobservation.net/2023/01/10/france-debates-its-near-space-policy/ (accessed on 11 August 2025).

5. Wilson P, Mantooth HA. Modeling Approaches. In: Model-Based Engineering for Complex Electronic Systems. Oxford, UK: Newnes; 2013.

6. Wilkinson PJ, Kelly TP. Functional hazard analysis for highly integrated aerospace systems. In: Proceedings of the IEE Certification of Ground/Air Systems Seminar (Ref. No. 1998/255). London, UK; 1998. p. 4/1-6.

7. S-18 Aircraft and Systems Development and Safety Assessment Committee. Guidelines and Methods for Conducting the Safety Assessment Process on Civil Airborne Systems and Equipment. Warrendale, PA, USA: SAE International; 2016.

8. Object Management Group. A UML Profile for MARTE, Beta 1; Technical Report ptc/07-08-04. Needham, MA, USA: Object Management Group; 2017.

9. Biggs G, Sakamoto T, Kotoku T. A profile and tool for modelling safety information with design information in SysML. Softw Syst Model. 2014;15:813-33.

10. Pereira DP, Hirata C, Nadjm-Tehrani S. A STAMP-based ontology approach to support safety and security analyses. J Inf Secur Appl. 2019;47:302-19.

11. Muller M, Roth M, Lindemann U. The hazard analysis Profile: Linking safety analysis and SysML. In: Proceedings of the 2016 Annual IEEE Systems Conference (SysCon). Orlando, FL, USA; 2016:1-7.

12. Yang L, Li S, Li C, Zhu C. Data-driven multivariate regression-based anomaly detection and recovery of unmanned aerial vehicle flight data. J Comput Des Eng. 2024;11:176-93.

13. Wang B, Liu D, Peng Y, Peng X. Multivariate regression-based fault detection and recovery of UAV flight data. IEEE Trans Instrum Meas. 2019;69:3527-37.

14. Jiang G, Nan P, Zhang J, Li Y, Li X. Robust spatio-temporal autoencoder for unsupervised anomaly detection of unmanned aerial vehicle with flight data. IEEE Trans Instrum Meas. 2024;73:3526014.

15. Ahn H, Chung S. Deep learning-based anomaly detection for individual drone vehicles performing swarm missions. Expert Syst Appl. 2024;244:122869.

16. Wang Y, Li S, Yang L, Zhang Y, Li C, Zhang A, et al. Deeply integrated autoencoder-based anomaly detection and critical parameter identification for unmanned aerial vehicle actuators. IEEE Sens J. 2024;24:24905-20.

17. Yang L, Li S, Zhu C, Zhang A, Liao Z. Spatio-temporal correlation-based multiple regression for anomaly detection and recovery of unmanned aerial vehicle flight data. Adv Eng Inform. 2024;60:102440.

18. Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, et al. Attention is all you need. In: Proceedings of the 31st Conference on Neural Information Processing Systems (NIPS 2017). Long Beach, CA, USA; 2017;30.

19. Gao Y, Hu X, Yang X, Xia X. Automated Unit Test Refactoring. In: Proceedings of the ACM on Software Engineering. New York, NY, USA: Association for Computing Machinery; 2025;2:713-33.

20. Altmayer Pizzorno J, Berger ED. CoverUp: Effective High Coverage Test Generation for Python. In: Proceedings of the ACM on Software Engineering. New York, NY, USA: Association for Computing Machinery; 2025;2:2897-919.

21. Lian X, Wang S, Zou H, Liu F, Wu J, Zhang L. Incorporating Verification Standards for Security Requirements Generation from Functional Specifications. In: Proceedings of the ACM on Software Engineering. New York, NY, USA: Association for Computing Machinery; 2025;2:1710-32.

22. Wei J, Wang X, Schuurmans D, Bosma M, Xia F, Chi E, et al. Chain of Thought Prompting Elicits Reasoning in Large Language Models. Adv Neural Inf Process Syst. 2022;35:24824-37.

23. Yao S, Yu D, Zhao J, Shafran I, Griffiths T, Cao Y, et al. Tree of thoughts: Deliberate problem solving with large language models. Adv Neural Inf Process Syst. 2023;36:11809-22.

24. Madaan A, Tandon N, Gupta P, Hallinan S, Gao L, Wiegreffe S, et al. Self-refine: Iterative refinement with self-feedback. Adv Neural Inf Process Syst. 2023;36:46534-94.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

25. Xie Y, Kawaguchi K, Zhao Y, Zhao JX, Kan MY, He J, et al. Self-evaluation guided beam search for reasoning. Adv Neural Inf Process Syst. 2023;36:41618-50.

26. ISO 26262. Road Vehicles—Functional Safety. Geneva, Switzerland: International Organization for Standardization; 2018.

© 2023 The Author(s). | Published by Agro Intelligence Research Press