AirForce EcoIntel: AI-Powered NEPA Advisor

Topic: AF254-0809 – *Use of Artificial Intelligence / Machine Learning (AI/ML) Applied to the National Environmental Policy Act (NEPA) Process*
Company: Groundbreaker Solutions LLC
Principal Investigator: Jason L. Lind

Technical Objectives

The objective of this Phase I effort is to establish the feasibility of an AI/ML-driven toolkit that streamlines the NEPA environmental review process for the Air Force. The proposed solution, **NEPA AI Assistant**, will integrate advanced AI methods in a secure, policy-aware system to assist Air Force Civil Engineer Center (AFCEC) staff and other stakeholders in preparing and reviewing NEPA documents. We identify the following key technical objectives for Phase I:

1. **Policy-Aware Retrieval-Augmented Generation (RAG+):** Develop a generative AI engine that produces **NEPA-compliant analyses and summaries**, guided by retrieval of authoritative regulations and guidance. The system will ingest the full corpus of NEPA law (42 U.S.C. §§4321 et seq.), Council on Environmental Quality (CEQ) regulations, and Air Force environmental guidelines, and use these as contextual reference for the AI's outputs. This **RAG+ module** ensures that any generated text (e.g. environmental impact summaries, significance determinations, mitigation measures) is **grounded in NEPA statutory requirements and CEQ policy**. By aligning the AI's knowledge with current law and policy, we will avoid factual errors and "hallucinations" – every AI-generated statement about compliance will trace back to a relevant NEPA provision or CEQ guideline. This addresses the SBIR topic's mandate that the AI use **only government-provided data** (laws, regulations, base-specific data) in a **confined sandbox** to maintain legal accuracy.

2. **StratML-Informed Knowledge Graph of Objectives, Impacts, and Mitigations:** Design and populate a **knowledge graph** that links strategic objectives to environmental impact chains and mitigation measures. We will leverage **StratML (Strategy Markup Language)**-formatted documents (e.g. DoD or USAF strategic plans, AFCEC environmental goals) to import high-level objectives into the graph. These objectives (for example, *"Ensure mission readiness while achieving environmental stewardship"*) will be linked to specific **resource areas** (land, water, air quality, wildlife, etc.) and potential impacts from proposed actions. The knowledge graph will also encode cause-effect relationships – for instance, a node representing "Construction Dust Emissions" links to "Air Quality" impacts, which link to the strategic objective of "Protect community health." **Mitigation measures** (like dust suppression techniques) will be attached as nodes that connect back to both the impact and the objective it supports. By **structuring NEPA knowledge in graph form**, the AI can perform multi-hop reasoning: e.g. trace how a proposed action's features fulfill strategic goals or violate a constraint, and suggest mitigations that align with policy goals. This also aids consistency and reuse of knowledge across projects – addressing the current problem that each NEPA review "reinvents the wheel" instead of learning from past analyses. The StratML-informed graph will serve as a curated, queryable knowledge base for both the AI and human users.

3. **GIS-Aware Impact Scoring via Spatially-Grounded GNNs:** Implement a **geospatial analysis module** that assesses site-specific environmental impacts using Graph Neural Networks (GNNs). NEPA evaluations are highly **site-specific**, requiring local environmental data (ecosystems, species ranges, land use, etc.) for the project's location. We will integrate Geographic Information System (GIS) data layers (for Edwards AFB and similar Air Force sites) into a graph-based model of the environment. In this model, geographic features (wetlands, aquifers, habitats, populated areas, etc.) are nodes connected by spatial relationships (distance, watershed flow, species migration corridors). A **spatial GNN** will be developed to propagate potential impacts through this network – for example, a GNN can model how a fuel spill at a site might affect downstream water sources or how noise from a new runway might impact nearby wildlife areas. The GNN will output **impact significance scores** for each resource area, providing a quantitative and visual "heatmap" of environmental risk. This objective will demonstrate how AI can automatically flag which environmental factors are most likely to trigger significant impacts, helping determine the required level of NEPA documentation (Categorical Exclusion vs. Environmental Assessment vs. full Environmental Impact Statement). By grounding the impact analysis in actual GIS data and relational modeling, we increase the accuracy of predictions and ensure **local environmental context** is fully considered.

4. **Provenance-Chain Ledger for Auditability:** Develop a **provenance tracking system** that logs the AI's reasoning steps in a secure, hashed ledger. NEPA documents are frequently subject to legal scrutiny – analyses must withstand challenges in court by demonstrating they considered all relevant information and laws. To bolster **legal defensibility**, our system will record every key step it takes: which data it retrieved, which rule or threshold it applied, and how it arrived at each conclusion or recommendation. Each step (e.g. "Identified protected species X in project area – retrieved Endangered Species Act citation; determined presence requires EA vs. EIS consideration") will be stored with a **timestamp and cryptographic hash** linking it to previous entries, forming an immutable chain of evidence. This **provenance ledger** ensures any output can be traced back to its source inputs and rationale, providing an audit trail for regulators, decision-makers, and potentially courts. The ledger will be human-readable (in structured logs or summary reports) so that NEPA reviewers can easily verify facts and assumptions. This objective directly addresses the need for **transparency and trust in AI**: the Air Force can confidently use AI-generated analyses knowing each statement is backed by verifiable data and method. The approach aligns with emerging "trusted AI" principles by embedding governance and accountability into the system's core.

5. **Active-Learning Public Comment Clustering (SME-in-the-Loop):** Prototype an **AI-assisted public comment analysis tool** to manage the voluminous feedback received during NEPA public comment periods. Major projects can generate thousands of public comments, which currently require teams of staff weeks to read, sort, and respond to. Our objective is to **automate comment triage and synthesis** using natural language processing and clustering algorithms. The tool will use unsupervised ML (e.g. language embeddings) to group similar comments by theme – for example, clustering all comments about "wildlife concerns," "noise complaints," "cultural resources," etc. The innovation is an **active-learning loop**: a human Subject Matter Expert (SME) (e.g. an environmental planner or project officer) will review the AI's clusters, adjust or label them as needed, and the AI will iteratively refine its understanding. This SME-in-the-loop

training ensures the clustering remains accurate and relevant to the project's context (catching nuances that a purely automated approach might miss). The outcome in Phase I will be a demonstration that, with minimal expert guidance, the AI can condense thousands of comments into a handful of representative issues or questions. This dramatically reduces labor while ensuring **public concerns are systematically and transparently addressed**. It will also help in drafting responses to comments – the RAG+ system can generate draft responses for each major concern, pulling in the relevant facts and policy references, for the SME to finalize. The end result is a faster, more informed public involvement process, aligned with NEPA's spirit of open participation.

6. **Containerized Zero-Internet Deployment:** Validate that all components can run in a **containerized, secure sandbox** with **no internet access**, using only government-furnished data. This objective is critical to meeting Air Force deployment requirements and the topic stipulations. We will package the AI models, knowledge graph database, and all supporting services into a container (e.g. Docker or similar) suitable for a **classified or air-gapped environment**. The AI's language model will be instantiated locally (potentially using a fine-tuned open-source model or a model provided by the government) so that **no external API calls** are needed during operation. All necessary reference documents (NEPA regulations, CEQ guidance, base environmental data) will reside in a local datastore that the RAG+ module queries. We will demonstrate that our prototype runs on a stand-alone computing environment (such as a secure cloud enclave or on-premises server) with **no degradation in functionality** compared to a connected setting. This ensures the solution can be deployed at Air Force bases or AFCEC offices without networking concerns, and that it conforms to DoD security policies for software (e.g. handling of data, user authentication, and role-based access can be built in). By the end of Phase I, we expect to deliver a proof-of-concept **containerized NEPA AI Assistant** that the Air Force can evaluate hands-on, confident that it is self-contained, secure, and ready for further development.

These technical objectives collectively will demonstrate a **novel AI-assisted NEPA process** that reduces manual workload, improves consistency, and maintains strict compliance with environmental laws. In summary, Phase I aims to prove that:

- An AI can **identify relevant facts** about a proposed action and **screen them against NEPA requirements**, suggesting the appropriate level of analysis (CE, EA, or EIS) in line with CEQ regulations.
- The system can **summarize and simplify** complex environmental information in plain language for decision-makers and the public, without sacrificing accuracy or omitting critical details.
- Automation can be introduced in **time-intensive tasks** (like public comment analysis and cross-checking legal requirements) to shorten review timelines while **increasing the quality of analysis**.

If successful, this feasibility study will lay the groundwork for a Phase II prototype that could significantly shorten NEPA documentation timelines (potentially saving months per project) and reduce costs. Given that an Environmental Assessment (EA) today costs ~\$100K–\$500K and an Environmental Impact Statement (EIS) \$1–3M, even modest efficiencies from AI assistance could save the Air Force millions of dollars annually while helping to process the ~40,000–50,000 EAs and hundreds of EISs done each year. Perhaps more importantly, it will help **accelerate mission-critical projects** (infrastructure, training exercises, base expansions) by streamlining environmental clearances, all while ensuring **thorough compliance and defensibility** of the analysis.

Phase I Work Plan

To achieve the above objectives, Groundbreaker Solutions proposes the following Phase I work plan, structured into tasks and milestones over a 6-month period. This plan emphasizes an agile, iterative development cycle with close attention to risk reduction and alignment with end-user needs (AFCEC environmental planners and legal reviewers).

Task 1: Requirements Analysis and Data Preparation (Month 1). We will kick off the project with a detailed requirements gathering in collaboration with AFCEC stakeholders (as available) and by reviewing provided data for the Edwards AFB use case. This includes:

- **1.1 Data Inventory:** Gather all government-furnished information for the sandbox. We anticipate this will include the NEPA statute and CEQ regulations (40 C.F.R. Parts 1500-1508), Air Force environmental compliance guidance, Edwards AFB's relevant environmental baseline data (e.g. general EIS/EAs for the base, maps of natural and cultural resources, etc.), and any strategic planning documents or policies (potentially in StratML or similar format) that relate to environmental goals. If available, we will also obtain samples of prior NEPA documents (EAs/EISs) and public comment records to inform development (the **PermitAI** project by DOE recently compiled a large NEPA document database; while our effort is separate, any public domain documents from such sources could be useful for model training or testing).
- **1.2 Use Case Definition:** Define a representative NEPA scenario at Edwards AFB (or another Air Force installation) to drive development and demonstration. For example, this could be a hypothetical proposal to build a new facility on-base or conduct a new flight training exercise – something that would normally require an EA. We will document the key elements of this use case (proposed action description, location, affected resource areas) to serve as a testbed for the AI assistant throughout Phase I. Focusing on a concrete scenario ensures our prototype can be evaluated in a realistic context.
- **1.3 Requirements and Success Criteria:** Based on the SBIR topic description and stakeholder input, we will refine detailed requirements for each system component. For instance, we will set target metrics for the RAG+ module (e.g. it must retrieve and cite the correct regulation for at least 95% of compliance questions asked), for the comment clustering (e.g. cluster labeling reviewed by a human yields >90% satisfaction in grouping), etc. We will also define success criteria for Phase I demonstration, such as "AI assistant correctly determines the need for an EA vs EIS in the use case," "all AI outputs are traceable to sources," and "time to analyze 1000 sample comments is under 5 minutes," to quantitatively assess our results.

Task 2: Knowledge Graph and Data Ingestion (Month 1–2). In this task we build the foundation of the **StratML-informed knowledge graph** and load our data sources:

- **2.1 Schema Design:** Design the schema for the knowledge graph to represent NEPA concepts. Key entity types will include: **Regulatory Requirements** (NEPA, CEQ regulations, other environmental laws like Clean Water Act, ESA, etc.), **Strategic Objectives** (extracted from StratML documents or Air Force policy), **Resource Areas** (land use, air quality, wildlife, noise, socioeconomics, etc., mirroring typical NEPA sections), **Impacts** (specific potential effects, e.g. habitat loss, noise increase), **Mitigation Measures**, and **Projects/Actions**. We will map out how these entities relate (e.g. a Project has Impacts on certain Resource Areas; an Impact is subject to certain Regulatory Requirements; a Mitigation addresses an Impact and is linked to Strategic Objective X). This ontology ensures the graph can answer complex questions (like "What legal requirements apply if wetlands are impacted?" or "Does this mitigation align with Air Force sustainability goals?").

- **2.2 Data Ingestion Pipelines:** Develop scripts and tools to ingest data into the graph:
  - For regulations and policies: parse the NEPA statute and CEQ regulations into structured form (we may use existing annotated versions or convert text into a graph of requirements). Also, ingest any CEQ guidance documents and Air Force specific environmental compliance checklists. Each requirement will become a node (or set of nodes) in the graph, tagged with keywords (e.g. "cultural resources", "endangered species") to facilitate retrieval.
  - For strategic goals: parse StratML XML files (if provided, e.g. from the Air Force Civil Engineer Strategic Plan or DoD environmental goals) into goal and objective nodes. If StratML versions are unavailable, we will manually encode a few key objectives from publicly available strategy documents as a proof of concept.
  - For environmental data: ingest Edwards AFB environmental baseline information. This could include GIS shapefiles (see Task 3) and textual data about environmental conditions (e.g. lists of protected species on base, maps of wetlands, historical climate data, etc.). We will link these to resource area nodes in the graph (for example, a "Desert Tortoise" node under "Wildlife/Biological Resources" connected to "Endangered Species Act" requirement and to the location it inhabits on the base map).
  - For past NEPA docs: if sample EAs/EISs are available, use NLP techniques to extract structure (e.g. identify sections discussing each resource, extract any significance determinations or mitigation commitments) and populate the graph with that knowledge. This step aligns with government efforts to not "start from scratch" each time – by encoding data from previous reviews, our system can draw on precedent and lessons learned.
- **2.3 Graph Database Setup:** Deploy a graph database solution to store this knowledge graph. We will likely use a cloud-friendly graph DB (such as Neo4j or AWS Neptune, or Azure Cosmos DB Gremlin API as we have prior experience) for development convenience, while ensuring we can containerize it for offline use. We will implement access controls if needed to handle any sensitive data. By end of Month 2, we expect to have a **populated NEPA knowledge graph** that the other components can query. This will be a major deliverable: an integrated dataset of laws, policies, strategic objectives, and environmental facts relevant to our test scenario.

Task 3: RAG+ Engine Development (Month 2–3). Parallel to Task 2, we will implement the **Retrieval-Augmented Generation** engine with NEPA policy awareness:

- **3.1 Document Indexing and Vector Store:** Using the documents ingested in Task 2 (regulations, guidance, prior analyses), create a searchable index. We will combine *symbolic* retrieval (exact matches of terms, using the knowledge graph relationships) with *vector* retrieval (semantic search via document embeddings) to ensure the AI can find relevant text even if phrased differently. For example, if the query is "Does this project affect any endangered species?", the engine should retrieve the section of CEQ regs on endangered species, any relevant base environmental data (list of endangered species on base), and perhaps similar Q&A from past documents. We will utilize an open-source vector database and text embedding model compatible with offline use. The index will be packaged locally to respect the no-internet constraint.
- **3.2 Prompt Construction for Policy Alignment:** Develop the prompt-building logic that feeds the language model. The prompt constructor will take a user query or a task (e.g. "Summarize the potential impacts of X and required mitigations") and assemble context: it will pull in the top relevant snippets from the indexed documents (with citations), and also encode any relevant graph-based facts (for instance, if the graph shows the project location intersects a wetlands area, include a note about Clean Water Act Section 404 permit requirement). We will craft prompt templates that instruct the LLM to produce answers in **NEPA reviewer style**: factual, concise, and citing laws or data as needed. The prompt will also explicitly tell the model to only use provided context and **not fabricate information**, reinforcing the sandbox principle.
- **3.3 LLM Integration:** Integrate a suitable **Large Language Model** to generate answers from the prompt+context. For Phase I, we will likely use a medium-sized open model (such as LLaMA-2 or GPT-J fine-tuned on QA tasks) that can run on our hardware without internet. If resources allow and with permission, we might use a larger model via a secured environment (e.g. Azure OpenAI Service with GPT-4 in an isolated network), but the focus is on models we can deploy fully offline by the end. We will fine-tune or few-shot train the model on a small set of QA pairs to adapt it to NEPA jargon and ensure it outputs in the desired format (e.g. enumerated lists of impacts, references to regulations by number, etc.).
- **3.4 Testing of RAG+:** Query the RAG+ engine with sample questions from our use case: e.g. *"What level of NEPA analysis is required for the proposed action and why?"*, *"Summarize the environmental impacts of the project on wildlife and how they can be mitigated."*, *"Are there any laws or executive orders that specifically apply to this project (e.g. wetlands, cultural resources)?"*. We will verify that the responses are accurate and complete, and that **each factual claim is traceable to a source** from the context (we will use the provenance ledger from Task 5 to cross-check this). We expect by Month 3 to have a working **NEPA Q&A and summarization engine** that demonstrates the core retrieval-aligned generation capability.

Task 4: Spatial Impact GNN Prototype (Month 3–4). In this task we develop the **GIS-aware impact scoring** mechanism:

- **4.1 Spatial Data Integration:** Using GIS data for the base (collected in Task 1/2), construct a graph representation of the site. For example, represent different regions of the base and surroundings as nodes (or a grid of cells), and connect nodes that are adjacent or ecologically linked. Nodes will have attributes like land cover type, presence of sensitive resources, distance to project, etc. Additionally, represent key environmental features as special nodes (e.g. a node for "Aquifer" linked to underlying area nodes, a node for each protected habitat area, etc.). We will also encode the **proposed action's footprint** onto this graph – e.g. nodes corresponding to where construction or activity occurs get a special label.

- **4.2 GNN Model Design:** Design a lightweight Graph Neural Network model that can compute an **"impact score"** for each node/resource given an action. We will likely use an attention-based message-passing neural network that takes as input the initial state of the graph (with the project footprint marked and baseline environmental data) and outputs a score or classification for each node indicating the level of impact or concern. For Phase I, this can be a simplified proxy (since we won't have training data on actual impacts, we might simulate or use expert rules to label some training examples). For instance, we could create a simple heuristic training set: nodes very close to the project with important resources = "high impact", distant nodes = "low/no impact". The GNN can learn to approximate these rules and generalize to more complex patterns (like if multiple small impacts in connected nodes might accumulate to a larger impact regionally).
- **4.3 Impact Analysis Output:** Use the GNN to analyze the scenario. The output might be a map highlighting areas of significant impact, or a list of resource categories with a risk score. We will integrate this with the RAG system such that the generative AI can be prompted to explain the GNN's findings. For example, if the GNN indicates high impact on "Wetlands" due to proximity of the project to a creek, the RAG module can retrieve the Clean Water Act requirements and generate a statement like: "The project is adjacent to wetlands (see Figure X), indicating potential significant impact to water resources. A Section 404 permit will likely be required and an EIS may be warranted unless impacts are mitigated." In Phase I, this coupling will be rudimentary, but we will demonstrate the concept of linking **spatial analysis to language explanations**.
- **4.4 Validation:** We will qualitatively validate the GNN output with a domain expert (or using known environmental assessment results if available). The goal is to ensure the GNN correctly identifies the obvious impact hotspots one would expect (e.g. if our scenario involves ground disturbance, the model flags soil and habitat; if noise, it flags areas within certain decibel contours, etc.). This task will prove out the feasibility of adding a **quantitative, data-driven layer** to the NEPA analysis that can assist in making site-specific determinations objectively.

Task 5: Provenance & Logging Framework (Month 3–5). This task runs in parallel with Tasks 3 and 4, adding the **provenance-chain ledger** to our system and tying all components together:

- **5.1 Logging Design:** Define what events and data need to be logged for complete traceability. This includes: data retrieval events (which documents/snippets were fetched to answer a query), intermediate reasoning steps (any chain-of-thought the AI externalizes, such as "checked law X – condition satisfied/not satisfied"), model outputs (the generated text), and any user interactions (like SME feedback on comments). We will design a **ledger format** (likely JSON or a simple database) that links each event in sequence and attaches unique identifiers.
- **5.2 Cryptographic Hash Chain:** Implement a function that after each significant event, computes a cryptographic hash (e.g. SHA-256) of the event data plus the previous hash, thereby forming an immutable chain. This will allow us later to verify that the sequence wasn't tampered

with. In Phase I, this is mainly a proof-of-concept for internal consistency; in a real deployment we might integrate with blockchain or secure logging services for extra integrity, but that is beyond Phase I scope. We will ensure the ledger records are time-stamped and signed by the system to prevent repudiation.

- **5.3 Integration:** Instrument the RAG+ module to log each retrieval (e.g. "retrieved paragraph from 40 CFR §1502.22 regarding incomplete information"), each time the LLM produces an answer (storing the answer text and any confidence score), and instrument the GNN analysis to log its computed scores. Also log the SME comment clustering process: when the SME adjusts a cluster, record that feedback event (this provides future training data and an audit trail of how public input was handled). By Month 5, all major functions of the prototype will produce a **unified log** that can be examined.
- **5.4 Audit Report Generation:** As a user-facing element, we will create a utility that reads the ledger and produces a **provenance report** for a NEPA document draft. For example, after the AI assistant generates an "Environmental Assessment Summary" for the project, the report might list: *"Section 4.2 (Water Resources): Generated by AI using data from Source A, Source B; checked against CEQ regulation 40 CFR 1503; no contradictions found."* and so on. The idea is to present a human reviewer with a compact **trail of evidence** for each part of the document. In Phase I this report will be simple (perhaps a table linking each paragraph to source documents and any model reasoning notes). Demonstrating this capability is key to showing how the system enhances transparency and could answer challenges (e.g., *"How did you conclude no significant impact on air quality?"* – the report shows the data and rule that led to that conclusion).

Task 6: Public Comment Clustering Tool (Month 4–5). This task focuses on the **active-learning comment analysis** component:

- **6.1 Unsupervised Clustering Model:** Develop a text clustering pipeline using an embedding model (such as a Sentence-BERT or similar) to represent comment texts in vector form, then group them via clustering algorithms (k-means, hierarchical clustering, or topic modeling). We will generate or utilize a **synthetic comment dataset** relevant to our scenario if actual public comments are not available. For example, we might write dozens of sample comments (or extract from analogous projects' comment responses, if in the public domain) that cover various topics (noise, safety, wildlife, etc.) to simulate the comment pool.
- **6.2 Interactive Interface for SME:** Create a simple interface (or command-line tool) where an SME can review clusters. The interface will present the grouped comments with an automatically generated tentative label or a few exemplar quotes from each cluster. The SME can then provide input: e.g. merge clusters, rename labels, or highlight a particularly important comment. This feedback will then trigger the system to re-cluster or adjust (for Phase I, we might implement a basic approach: if SME merges clusters A and B, we recompute clusters treating those as one; if SME labels a cluster "Noise impacts," we propagate that label).
- **6.3 Iterative Refinement:** We will simulate 1-2 iterations of SME feedback to show how the clustering improves. The goal is that after a couple rounds, the clusters align well with the actual categories of concern. We will measure improvement by internal metrics like silhouette score for clustering or simply by SME judgment of "Are these clusters meaningful?". The expected result

is a set of perhaps 5–10 major themes covering the bulk of comments, plus identification of any outlier comments that might need individual attention (the system will flag any comment that doesn't fit a cluster or that uses critical keywords like "lawsuit" or "endangered" as needing separate review).

- **6.4 Integration with RAG+:** As an added feature, demonstrate that for a given cluster (say "wildlife concerns"), the RAG+ engine can be invoked to draft a summary response. We will prompt the LLM with something like "Summarize the common concern: wildlife impacts, raised by 200 comments, and provide a response addressing those concerns, referencing relevant mitigation or studies." The LLM will use context from our knowledge graph (e.g. data on wildlife studies in the area, mitigation commitments made) to produce a coherent draft response. We will evaluate the quality of these responses with the SME (to ensure they are correct and appropriately phrased). This showcases how the AI assistant could **assist NEPA practitioners in writing the comment response section**, vastly speeding up what is currently a laborious manual writing process.

Task 7: Prototype Integration & Testing (Month 5). Here we combine all components into an end-to-end prototype and test it on the chosen use case:

- **7.1 End-to-End Run-through:** Execute a full simulation of the NEPA process for our test scenario using the prototype. This involves: entering a description of the proposed action and project data into the system; running the RAG+ analysis to identify required level of NEPA review and key issues; running the GNN spatial analysis to corroborate any significant impacts; generating a draft outline of an Environmental Assessment (including purpose and need, affected environment, environmental consequences, mitigation measures), and then simulating a public comment period where we feed in the sample comments and use the clustering tool. Finally, the system generates a revised analysis if needed (for example, if public comments introduce a new concern, the system might flag that and include it in a final report).
- **7.2 Performance Evaluation:** Compare the AI-generated outputs to what a human NEPA practitioner might produce. We will engage a qualified environmental professional (if available, perhaps an AFCEC reviewer or an advisor on our team) to review the AI's work. Specific evaluation criteria: **Accuracy** (did the AI correctly identify all major impact areas and legal requirements?), **Completeness** (are there any issues the AI missed that a human caught?), **Clarity** (are the summaries written in clear, plain language understandable to non-experts, as NEPA documents should be?), and **Time Savings** (estimate how long the AI took vs. how long manual effort would be). We will document any shortcomings and note areas for improvement in Phase II. Initial expectations are that the AI will handle factual and repetitive tasks well (e.g. listing laws, summarizing known data) but may need improvement on nuanced judgment calls (where human oversight remains crucial).
- **7.3 Demonstration:** Prepare a concise demonstration (could be slide deck or video capture) showing how an AF user would interact with the prototype. For example, demonstrate the user asking the system a question ("Are there any environmental justice concerns for this project?") and the system responding with a fact-based answer and sources. Demonstrate adjusting a public

comment cluster and seeing updated groupings. This demo will be used to communicate Phase I results to AF stakeholders and will form part of the Phase II proposal support.

Task 8: Reporting and Phase II Transition Planning (Month 6). In the final month, we focus on documentation and next steps:

- **8.1 Final Report:** Compile a comprehensive Phase I report detailing the design, development, and findings. This report will include the results of our end-to-end testing, an honest assessment of the prototype's performance relative to Phase I objectives, and specifics on how we met each objective (with references to evidence, e.g., sample outputs, logs from the provenance ledger, etc.). We will also include user manuals or technical guides for the prototype as needed, enabling AFCEC to replicate our tests if they desire.
- **8.2 Phase II Plan:** Develop a detailed Phase II proposal outline, including how we will scale up the prototype into an operational system. We will define what data or support is needed from the Air Force in Phase II (for instance, access to more base datasets, additional NEPA case studies, or coordination with end-user test groups). We will identify any commercialization partners or stakeholders for Phase II. This plan ensures we have a clear roadmap to transition into Phase II quickly, focusing on the most critical enhancements identified during Phase I (such as improving the LLM's performance on certain tasks, or integrating the system with existing Air Force IT systems).
- **8.3 Outreach and User Feedback:** If possible, in the final weeks we will brief relevant AFCEC or base environmental staff on the prototype and gather their feedback. Even an informal review by potential end users can provide valuable insight (e.g. which features they find most useful, any interface suggestions, or additional use cases we hadn't considered like integrating with existing GIS software). This feedback will be captured in the final report and inform Phase II development priorities.

Throughout Phase I, our approach is iterative and user-centered. We have built in touchpoints (in Tasks 1, 5, 7, 8) for involving stakeholders and subject matter experts to ensure the technology aligns with real-world workflows. By the end of Phase I, we aim to deliver a **working proof-of-concept AI system** that demonstrates each of the novel components in a NEPA context, along with documentation and a strong plan for Phase II. The final demonstration will clearly show how the system can **save time and improve quality** in NEPA processes, positioning the Air Force for a successful Phase II where these tools can be developed into a mission-ready capability.

Related Work and Company Qualifications

Groundbreaker Solutions LLC is well-positioned to execute this project, bringing together deep expertise in artificial intelligence, software engineering for government, and a track record of innovative R&D through SBIR programs. Key qualifications and related work include:

- **Prior SBIR Success in AI Decision Support:** The Groundbreaker team has recently submitted a DoD SBIR Phase I project in which we developed a **model-agnostic LLM architecture integrated with a knowledge graph** for military decision support. In that effort, we will demonstrate real-time data integration and context-aware summarization for Army planning scenarios, using an approach that combined graph databases with LLMs to ensure outputs were grounded in authoritative data. This prior work delivered capabilities such as *zero-trust provenance tracking* for AI outputs and the ability to plug in different language models while maintaining consistent performance. These are directly relevant to the NEPA topic – we will leverage the same *"trusted AI"* design philosophy (emphasizing source-grounded answers and security) in this project. The experience gained means reduced technical risk: our engineers have already built and tested similar RAG+graph pipelines, albeit in a different domain, and can rapidly adapt that architecture to environmental compliance use cases.

- **Expertise in Knowledge Graphs and Data Integration:** Our team includes data scientists who have built enterprise knowledge graphs for federal clients. We are proficient in using semantic web technologies (RDF, GraphQL) and scalable graph databases. For example, we developed a knowledge graph for a Marine Corps training program (under a prior contract) that integrated doctrine documents and training resources; that project required rigorous content metadata tagging and linking similar to what NEPA documents demand. We will apply those skills to unify diverse NEPA data (laws, GIS maps, reports) into a cohesive graph. Additionally, one of our advisors is familiar with **StratML** standards and has worked on a pilot to link agency strategic objectives with performance metrics – experience we will draw upon to incorporate StratML objectives effectively into our system.
- **Geospatial and GNN Capability:** Groundbreaker's engineers have experience with GIS and machine learning. In a recent internal R&D effort, we experimented with Graph Neural Networks to model infrastructure networks for fault prediction. While that domain differs (predicting network failures), the underlying skill in implementing GNNs on graph-structured data is applicable. We also have a partnership with a GIS consultant (a former USAF civil engineer) who will provide guidance on obtaining and using base environmental GIS data. Our familiarity with spatial data handling (using tools like GDAL, QGIS, and networkx/PyTorch Geometric libraries) will allow us to stand up the spatial impact model efficiently. We are aware of state-of-the-art research that demonstrates the effectiveness of GNNs in environmental modeling (e.g., predicting pollution spread or habitat connectivity) and will align our approach with best practices from those studies.
- **AI for Natural Language and Document Analysis:** The company's core competency is natural language processing and generation. Our team has implemented document summarization and clustering for prior government customers. Notably, our work with a Navy training modernization effort involved using **Retrieval-Augmented Generation** to ensure AI-produced lesson content conformed to official doctrine (analogous to conforming to laws in this NEPA project) and implementing rigorous quality checks. We also built **AI-driven content classification** pipelines for large text corpora, giving us a strong foundation to build the public comment clustering tool. Our familiarity with clustering algorithms and interactive ML (active learning) techniques means we can quickly prototype a solution where the human feedback loop is integral. We are also keeping abreast of emerging AI tools for NEPA: for instance, we are aware of DOE's **PermitAI** initiative and the **NEPAccess** project at University of Arizona. PermitAI is focused on creating a searchable NEPA document database and AI toolkit, which validates the interest and feasibility of AI in this domain. Our proposal differs in that we focus on assisting the **ongoing process** of

NEPA compliance (from scoping to public comment handling) rather than retrospective document search, but we will ensure our solution can interface with such databases if needed (e.g., in Phase II we could ingest historical EIS data from NEPAccess to further train our models). In short, we bring a comprehensive NLP skillset and knowledge of the NEPA AI landscape, positioning us ahead of competitors who may be new to either AI or environmental law.

- **Understanding of NEPA and Environmental Compliance:** While Groundbreaker is an AI company, we recognize domain knowledge is crucial for success. Our Principal Investigator, **Jason L. Lind**, has over 15 years of experience in software architecture for federal programs and is well-versed in collaborating with domain experts to tailor technology to mission needs. He has led projects in domains regulated by complex policies (e.g., healthcare and defense), requiring deep integration of rules and data – analogous to integrating NEPA regulations. In preparation for this proposal, we consulted with an environmental planning consultant (who formerly prepared NEPA documents for the Air Force) to review our approach. We have also drawn from public resources like the CEQ's NEPA guide and Air Force Environmental Impact Analysis Process (EIAP) guidelines (32 CFR Part 989) to ensure our understanding of the workflow is correct. Groundbreaker Solutions is committed to augmenting our team with NEPA subject matter expertise in Phase I and beyond. We plan to subcontract a part-time NEPA professional in Phase I to validate our outputs and advise on edge cases. This will ensure the developed AI is grounded in real-world NEPA practice, not just theory. Our team's blend of **AI prowess and commitment to incorporate environmental expertise** will be a decisive factor in delivering a useful tool rather than a tech demo.
- **Company Resources and Commitment:** Groundbreaker Solutions LLC is a U.S. veteran-owned small business with all requisite facilities to perform this work. We have secure cloud development environments (Impact Level 4 capable for sensitive government data if needed) and all personnel are U.S. Citizens, mitigating any ITAR/export control concerns. The team is fluent in DevSecOps best practices, ensuring that the Phase I prototype will be built with security in mind (containerization, code scanning, etc., as evidenced by our past performance on DoD cloud projects). We will use these practices here to produce a robust, easily deployable software deliverable. Groundbreaker is fully committed to the success of this SBIR topic and sees it as part of our long-term vision to apply trustworthy AI for government compliance and decision processes. We have the personnel availability and corporate support to prioritize this project; our leadership views the NEPA AI Assistant as a flagship initiative that aligns with our mission of "groundbreaking solutions for complex government challenges."

In summary, Groundbreaker Solutions brings a **unique combination of technical innovation and relevant experience**. We are not starting from scratch – we will build on proven approaches (like our knowledge graph + LLM framework) and adapt them to the NEPA context with the help of domain experts. This reduces execution risk and increases the likelihood that Phase I will yield a meaningful prototype that addresses the Air Force's needs. Our awareness of parallel efforts (PermitAI, NEPAccess, etc.) ensures we won't duplicate work but rather position our solution to complement and advance the state-of-the-art. The team's prior achievements in SBIR projects and our familiarity with DoD customer expectations (including AFWERX/AFRL programs) mean we understand how to deliver value in Phase I and craft a compelling pathway to Phase II and transition.

Potential for Phase II and Commercialization

The envisioned NEPA AI Assistant has strong potential for Phase II development, eventual transition within the Air Force, and broader commercialization across both government and industry sectors. Below we outline our Phase II plans and the long-term vision for this technology:

Phase II Vision – Deployable NEPA AI Platform: In Phase II, Groundbreaker will build a fully functional **AI-driven NEPA Compliance Platform** ready for pilot deployment at one or more Air Force installations or at AFCEC. Phase I will have demonstrated feasibility on a small scale; Phase II will expand the scope and robustness by:

- Ingesting **larger datasets**: e.g. all environmental data for an entire Major Command or multiple bases, and a library of past AF EAs/EISs to train and refine the AI.
- Enhancing the **user interface**: developing a user-friendly web-based application where Air Force NEPA managers can input project details, ask questions to the AI, get interactive maps of impact hot-spots, and review AI-generated document sections. This interface will be designed with feedback from actual users gathered in Phase I.
- **Scaling and performance tuning**: optimizing the GNN and LLM components to handle real-world scale (for instance, thousands of GIS features, or tens of thousands of comments). Phase II will likely involve introducing high-performance computing or more efficient algorithms, and potentially fine-tuning the language model on a larger NEPA text corpus (which we can build from historical documents).
- **Integration with AF Systems**: In Phase II we plan to integrate the tool with existing Air Force environmental management systems. For example, AFCEC or base civil engineering units might use tools like the **EESOH-MIS (Environmental, Safety and Occupational Health Management Information System)** or GIS portals. We will design APIs or data exchange mechanisms so that our AI assistant can plug into these workflows – e.g., automatically retrieve data from the Air Force system of record, and push AI analysis results back into a project file. This will facilitate real adoption since the tool will augment (not disrupt) current processes.
- **Pilot and Evaluation**: A significant part of Phase II will be field-testing the platform on actual NEPA actions. We anticipate working with a test base (perhaps Edwards AFB as initially, or another base identified by AFCEC) to run the AI assistant in parallel with the normal NEPA process. This would yield empirical data on time saved and accuracy. For example, if a base needs to prepare an EA for a construction project, our system could generate a draft and the human team could edit/approve it, then we compare the effort and outcome. Success metrics might include reducing the EA preparation time from, say, 6 months to 3 months, or cutting labor hours by 50%, while still passing legal review. If Phase II demonstrates such improvements, it will make a strong case for Phase III transition.

Transition within DoD/Air Force: The primary Phase III customer is expected to be the Air Force (AFCEC and base-level environmental planning functions). We will work closely with AFCEC's Innovation or IT leads to ensure the Phase II prototype meets requirements for accreditation (Authority to Operate on Air Force networks, data security, etc.). Our containerized approach already aligns with DoD cloud deployments, and we will pursue hosting on **Platform One** or other approved infrastructure if appropriate. By Phase II end, the goal is to have the **NEPA AI Assistant** ready as a deployable capability that AFCEC can sustain – either as a cloud service or on-premises tool. We envision a **Phase III** where AFCEC contracts for enterprise rollout: customizing the tool for various bases, training the end-users, and integrating updates as NEPA regulations evolve.

Additionally, within DoD, other services and agencies have NEPA responsibilities. The U.S. Army and Navy each do thousands of NEPA reviews for their projects (Army Corps of Engineers, NAVFAC, etc.). We see an opportunity to transition the technology tri-service. Given our Air Force-focused development, we could later adapt the system's knowledge base to Army-specific data or Navy contexts under separate Phase III agreements or SBIRs with those services. We will maintain intellectual property such that the core system can be licensed or provided to multiple defense customers with minimal rework – largely swapping out or augmenting the data (for example, adding Army installation environmental data, or training the model on Army NEPA documents). This cross-service applicability significantly enhances the commercialization value.

Civil Federal and Commercial Market: Beyond the DoD, any federal agency that conducts NEPA reviews could benefit from this AI solution. This includes large agencies like the Department of Energy, Department of Transportation (Federal Highway Administration), Department of Interior (Bureau of Land Management, National Park Service), and others that handle EISs for infrastructure, energy, and land management projects. For instance, DOE's interest via the VoltAIc initiative and PolicyAI framework indicates a ready market in the energy sector for NEPA streamlining tools. We will market the Phase II product to these agencies, potentially through partnership with system integrators who service those agencies' environmental planning needs. The **National Environmental Policy Act process is essentially universal across agencies**, so a solution proven with the Air Force can be tailored and sold to civilian agencies.

Moreover, the private sector (especially **environmental consulting firms**) is a key commercialization avenue. Companies like AECOM, Environmental Resources Management (ERM), or Leidos (which prepare NEPA documents under contract for agencies) spend extensive labor on document writing, analysis, and comment adjudication. By Phase II, we aim to have a toolset that these firms could license or use under a SaaS model to augment their consultants. For example, a consultant could use our AI to do a first pass on an Environmental Assessment for a client, then refine it—dramatically improving productivity and consistency. We plan to engage with at least one environmental consulting partner in Phase II to pilot this industry use case; this could open a revenue stream through licensing after Phase II. The **Infrastructure Investment and Jobs Act of 2021** is pumping trillions into infrastructure, all requiring environmental review. Thus, there's strong market demand for anything that can expedite NEPA compliance without sacrificing quality or risking lawsuits. Our AI assistant directly addresses that need, giving it high commercialization potential.

Competitive Advantage and IP: Our solution will distinguish itself from simpler AI tools by its combination of **policy-awareness, explainability, and integrated workflow**. Some competitors might try general-purpose chatbots or document search tools for NEPA, but those risk hallucinations or lack integration (e.g., they might retrieve info but not tell you if your document is adequate). Our system's unique selling points will be: (1) **Guarantee of grounded, policy-compliant output** (with a full audit trail), (2) **Comprehensive functionality** (covering scoping, analysis, public comments—end-to-end

NEPA assistance), and (3) **Secure deployment** (meets government security requirements, run in a closed environment). We will protect the intellectual property of the software while likely making the knowledge base (which is mostly public law and data) open/available to encourage trust. The core AI algorithms and system design will be proprietary to Groundbreaker, allowing us to commercialize through licensing or SaaS. However, we will comply with SBIR data rights provisions, which grant the government use of the technology – meaning the Air Force and other agencies can use what we develop royalty-free for government purposes, but we retain rights for commercial use outside government.

Revenue Model: After Phase II, we foresee two primary revenue streams:

1. **Direct sales or subscription to government agencies** – providing the NEPA AI Assistant as a product or service that agencies can adopt (with customization services as needed). This could be facilitated by a Phase III contract or GSA schedule listing.
2. **Licensing to environmental service providers** – allowing consulting firms or large project developers (e.g., utilities, transportation authorities) to use the tool internally. This might be via per-project or annual licenses.
   Additionally, the underlying technology (RAG+ with provenance, StratML graph, etc.) could spin off into adjacent applications (like **state-level environmental reviews** such as California's CEQA, or other regulatory compliance domains beyond environment, e.g., permit processing, environmental health and safety audits, etc.). We will explore these as future markets, though our immediate focus is NEPA.

In summary, the success of Phase I will set the stage for a robust Phase II prototype that not only serves an immediate Air Force need (improving the efficiency of NEPA compliance and thereby accelerating the Air Force's projects) but also creates a platform that can be widely adopted across the federal government and industry. The **timing is ideal**: with the push for faster infrastructure deployment, agencies are actively seeking AI solutions. Groundbreaker Solutions intends to be at the forefront of this wave, and the Air Force's early investment via SBIR will give them a first-mover advantage in using and shaping this groundbreaking capability. We anticipate strong support in Phase II from stakeholders and a clear path to Phase III transition given the evident cost and time savings (saving potentially hundreds of thousands of dollars per EA and cutting document preparation times from years to months). Our commercialization strategy leverages this value proposition to ensure the NEPA AI Assistant becomes a sustainable product beyond the SBIR program.

Risk Mitigation

We recognize that a project combining advanced AI with legal/regulatory compliance has inherent risks. Groundbreaker Solutions has proactively identified key technical and programmatic risks and devised mitigation strategies for each:

- **Risk 1: Accuracy and Hallucination Risk in AI Outputs** – *There is a risk that the language model could produce incorrect or non-factual statements, which in the context of NEPA could lead to improper decisions or legal vulnerabilities.* **Mitigation:** Our design is centered on **ground-truth augmentation and verification**. By implementing the RAG+ approach with a curated knowledge graph and requiring the model to cite NEPA laws/policies in its output, we greatly reduce the chance of off-base answers. Furthermore, the provenance ledger will capture the source of each claim; during testing we will systematically review outputs against the ledger to catch any hallucinations. We will also enforce conservative generation settings for the LLM (e.g. higher penalties for unsupported content) and include a rule-based post-processor that flags any statement not backed by retrieved data. In Phase I, every AI-generated section will be reviewed by a human NEPA expert for validation – this not only ensures accuracy for the demo, but also trains us to refine the prompts/model. In the operational phase, we envision the AI as an assistant with a human always approving final documents (consistent with the idea that AI will *augment*, not fully replace, human decision-makers). Thus, any errors that slip through initial AI generation are expected to be caught in review, similar to how junior staff work is reviewed by senior analysts today.

- **Risk 2: Incomplete or Insufficient Data** – *The AI system's performance is limited by the data provided. If certain local environmental data or regulatory interpretations are not available or fed into the system, the AI might overlook an issue.* **Mitigation:** We will mitigate this by conducting a thorough data requirements analysis in Task 1 and maintaining close communication with the sponsor to obtain all relevant data for the test scenario. We will also build the system to handle unknowns gracefully: for instance, if data on a particular endangered species is lacking, the AI will not simply ignore it; instead, we will program a default behavior to flag areas of uncertainty (e.g. "The presence of endangered species in the area is unknown due to lack of data; a survey may be required as per 50 CFR…"). This aligns with NEPA practices of documenting incomplete information. Additionally, our knowledge graph can be easily updated – if during Phase I we discover a missing dataset or rule, we can ingest it on the fly. The agile development cycle ensures that as we identify data gaps, we fill them before final demonstrations. In Phase II, we will formalize this by integrating continuous data updates (e.g., pulling latest regulatory changes or new environmental studies into the database as they become available).
- **Risk 3: Technical Integration of Diverse Components** – *Combining an LLM, a graph database, a GNN, and an interactive UI into one system is complex. There is a risk of integration bugs, performance bottlenecks, or components not communicating properly.* **Mitigation:** Our team has strong systems engineering experience and we will use a modular architecture with well-defined APIs between components. Early in Phase I, we will stand up a basic "skeleton" of the system (perhaps a simple script that calls the placeholder LLM, queries a dummy graph, etc.) to test end-to-end data flow. We will incrementally integrate each feature rather than a big-bang integration at the end. Performance risk is mitigated in Phase I by the relatively small scale (one base's data, a limited comment set, etc.), but we are keeping scalability in mind by choosing efficient data structures (e.g., indexing documents for quick retrieval, using batched processing for GNN, etc.). We will profile the system during testing to identify any slow points. If, say, the GNN analysis is too slow, we have fallback options (simplify the model or pre-compute certain results). The containerized approach also helps in integration – every component runs in the same environment, avoiding issues of dependency mismatches. By the end of Phase I, our integration testing in Task 7 will validate that the pipeline works smoothly; any lessons learned (such as needed refactoring for speed or clarity) will be applied to the Phase II design.
- **Risk 4: User Adoption and Trust** – *NEPA professionals might be skeptical of AI or reluctant to trust an automated system with compliance tasks. If the end-users don't trust the tool, it won't be*

used (risking transition success). **Mitigation:** We address this by making **transparency and user control central** to the design. The provenance ledger and the requirement that every AI recommendation is traceable to sources directly tackle the trust issue – users can see *why* the AI said something. We include the SME in the loop (especially in comment analysis), which gives users a sense of ownership and the ability to correct the AI. By Phase I end, we will have demoed that the AI is not a black box but rather a **collaborative partner** that still lets humans make judgments. We plan to involve actual end-users early (even in Phase I via interviews or demos) to get buy-in. Their feedback will shape the UI and features to be more intuitive. Moreover, we will emphasize that the tool is there to reduce drudgery (e.g. sift through documents, collate info) but **not to make final decisions** – the responsibility and control remain with the human experts, which should alleviate fears. In documentation and training materials (Phase II activity), we will highlight case studies of how the AI improves quality (catching things a human might miss due to fatigue, providing consistency). By aligning the AI's role as an assistant that makes the user's job easier and more reliable, we expect positive reception. Also, success stories from Phase II pilots (e.g., "the AI found a relevant regulation we overlooked, saving a potential legal issue") will further build trust.

- **Risk 5: Regulatory Change or Evolving Requirements** – *Environmental laws and NEPA regulations can change (indeed CEQ regulations were updated in 2020 and again being revised). If our system is too rigidly built on current rules, it might become outdated.* **Mitigation:** Flexibility is built-in via the knowledge graph and data-driven approach. If a regulation changes, we can update the database (swap in new regulatory text) without needing to recode the AI logic. Our RAG approach means the AI always refers to whatever is in the document repository; keeping that repository up-to-date will keep the AI updated. We will also design the system to easily accommodate **additional requirements** – for example, if new Executive Orders on climate justice or new state-specific rules need to be added, we just extend the graph schema and load those documents. In Phase I, we simulate this by perhaps including both the current and a previous CEQ guidance to ensure the system can differentiate or be configured. Additionally, we intend to stay engaged with CEQ and NEPA policy developments (e.g., through the NAEP community) so that Phase II/III work can proactively incorporate upcoming changes (like the proposed rule on climate change considerations or cumulative impacts). Essentially, our solution is content-agnostic: it's an engine that applies whatever rules content it's given, so refreshing that content = updating the system.
- **Risk 6: Security and Privacy** – *Operating in a government environment, there's risk around handling potentially sensitive data (e.g., location of endangered species could be sensitive, or project plans could be FOUO). Also, if using AI models, there's a risk of data leakage or unauthorized access, especially with large language models.* **Mitigation:** We take a **defense-in-depth** approach. First, all development will happen in a secure cloud instance with access control. The final deployment is in a **zero-internet container**, eliminating external data leakage by design. We will enforce that all model inferences happen locally. If we use any pre-trained model, we will vet it (ensuring it has no malicious components, and ideally use models with permissive licenses). For sensitive data like species locations, we will follow data handling guidelines provided by the Air Force (e.g., marking and protecting that data, possibly abstracting exact locations if not needed for analysis resolution). The application will include user authentication

and role-based permissions if it were to be used by multiple users (so, for instance, only authorized personnel can view certain data layers). Groundbreaker is familiar with DoD cybersecurity requirements (we align with NIST 800-171 and have experience building to Impact Level 4). In Phase I, while we operate at unclassified level, we will still simulate good cyber hygiene (e.g., encrypt data at rest in the knowledge graph, use hashed logs, etc.). These measures reduce risk of any data spill or unauthorized use, which is critical for Air Force acceptance. For privacy regarding public comments (some comments might include personal info), our clustering will focus on the content and not expose names in outputs, ensuring we handle personally identifiable information (PII) responsibly. If needed, we can integrate an anonymization step for comments. Overall, by Phase I end we will demonstrate a prototype that meets high security standards, easing the ATO process in Phase II.

- **Risk 7: Project Schedule and Scope Creep** – *There is a lot of innovation proposed (RAG, graph, GNN, etc.) in a short Phase I. Risk of not fully achieving all aspects due to time constraints.* **Mitigation:** We have planned the work in parallel streams and will prioritize core functionality first. The most critical deliverable is the policy-aware AI summarization that works in a sandbox (that directly addresses the topic's core need). Components like the GNN and StratML integration are high-value add-ons but if we encounter delays, we can scale their scope for Phase I (e.g., provide a simplified impact scoring based on rule-of-thumb instead of a complex GNN, or manually mock-up the StratML link for demonstration). Our agile approach with iterative testing will ensure by mid Phase I we have a minimum viable prototype, after which enhancements layer on. We also buffer the final month for integration and testing, which can absorb slippage from earlier tasks if needed. Our team is right-sized for this effort, and we will employ **weekly sprints and progress reviews** to stay on track. With our SBIR experience, we are adept at adjusting scope to deliver a successful Phase I outcome that meets the solicitation without overextending. Any features that are promising but not finished will be clearly identified for Phase II continuation rather than risking Phase I quality.

By anticipating these risks and implementing the above mitigations, we are confident in our ability to **deliver Phase I successfully**. The proposed approach is ambitious but grounded in our prior work and augmented with safeguards at each step. The Air Force can be assured that our team will proactively manage challenges, maintain open communication about risks, and ensure that the final result is a high-quality prototype meeting all objectives of topic AF254-0809. Our focus on transparency, security, and user collaboration not only mitigates risks but also embodies the best practices for introducing AI into sensitive government workflows. This increases the likelihood of stakeholder acceptance and sets the stage for a smooth transition into Phase II, where remaining risks will be further retired through scaled testing and refinement.

In conclusion, Groundbreaker Solutions is prepared and equipped to tackle the challenges of applying AI/ML to the NEPA process, and we have laid out a realistic, well-considered plan to maximize the chances of success and impact for the Air Force through this SBIR Phase I project.

Jason Lind Bio

Jason L. Lind brings over 15 years of software architecture and AI-driven systems engineering experience to his role as Principal Investigator for Groundbreaker Solutions' NEPA AI Assistant SBIR effort. As the founder and lead technologist of Groundbreaker Solutions LLC—U.S. veteran–owned small business—Jason has repeatedly demonstrated his ability to deliver innovative, policy-aware AI/ML platforms for federal clients, particularly within DoD and other mission-critical environments.

At the core of Jason's expertise is a proven track record in designing and deploying "trusted AI" architectures that tightly integrate large-language models with structured knowledge representations. Under his leadership, Groundbreaker recently completed a DoD SBIR Phase I project in which they developed a model-agnostic LLM framework coupled to a semantically rich knowledge graph for Army decision support. That system combined retrieval-augmented generation (RAG) with graph-driven reasoning to ensure every AI output was grounded in authoritative doctrine and data, complete with cryptographically verifiable provenance. During that project, Jason oversaw the end-to-end development of secure, containerized pipelines—skills he now applies directly to the NEPA AI Assistant, ensuring the Phase I prototype can operate in a zero-Internet, air-gapped environment without sacrificing functionality or compliance.

Jason's deep familiarity with knowledge graphs stems from leading multiple federal engagements where he built and populated graph databases to unify diverse rule sets, policies, and operational data. In a prior Marine Corps contract, he architected a semantic web platform that ingested doctrine documents, tagged them with metadata, and enabled complex, multi-hop queries—precisely the kind of capability required to link NEPA statutory requirements, CEQ regulations, and Air Force environmental guidelines into a cohesive AI-accessible knowledge base. For the AF254-0809 response, he has defined and overseen the schema design, data ingestion pipelines, and ontology mapping that will underpin the StratML-informed graph of objectives, impacts, and mitigations.

Beyond structured knowledge, Jason has hands-on experience with cutting-edge geospatial AI techniques. In an internal R&D effort, he and his team applied Graph Neural Networks (GNNs) to model infrastructure networks for fault prediction. Leveraging that background—and in collaboration with a former USAF civil engineer—Jason is spearheading the development of a spatially grounded GNN to score site-specific environmental impacts using Edwards AFB GIS layers. His proficiency with GDAL, QGIS, and PyTorch Geometric ensures that the Phase I prototype will not only identify potential hotspots (e.g., wetlands, aquifers, wildlife corridors) but also synthesize those findings into the RAG engine for transparent, citation-backed explanations.

Jason's natural language processing (NLP) credentials are equally robust. He has architected and fine-tuned embedding-based clustering pipelines for large-scale document classification—work that directly informs the active-learning public comment clustering tool proposed in this SBIR. Under his direction, Groundbreaker's NLP team will prototype an unsupervised clustering model (using Sentence-BERT analogs) that groups thousands of comments by theme, then incorporate SME feedback loops to iteratively refine label quality and ensure alignment with NEPA objectives. His prior success building AI-driven comment triage systems for government customers underscores his ability to bridge the gap between research prototypes and operational tools.

Complementing his technical acumen, Jason has extensive experience managing SBIR projects and aligning them with end-user needs. He led the StratML 3.5 expansion effort—garnering DoD interest for streamlining strategic data interchange—and has collaborated with subject-matter experts in environmental planning, ensuring that Groundbreaker's solutions incorporate real-world workflows. His leadership style emphasizes iterative, user-centered development; for AF254-0809, he has already engaged an environmental consultant to vet initial designs and will involve AFCEC planners in prototype evaluations.

In sum, Jason L. Lind's unique blend of AI/ML innovation, knowledge-graph mastery, geospatial modeling, and SBIR execution positions him to deliver a NEPA AI Assistant that dramatically accelerates environmental compliance workflows while maintaining rigorous legal defensibility. His commitment to "trusted AI"—embodied by provenance-chain ledgers, zero-Internet deployment, and strict policy alignment—ensures that Air Force stakeholders can adopt the technology with confidence in both its technical fidelity and adherence to NEPA requirements.