PROJECT #7: Fine-tuning specialized AI models on classified nuclear physics datasets for stockpile stewardship

STAGE 1: UCDP v9.3-SFE VISION GENERATION

Input:
Main Idea: Fine-tune specialized AI models on classified nuclear physics datasets for stockpile stewardship using AI.
Run Mode: Standard (Full Protocol)
Sector: Defense/Nuclear Security

1. **Cross-Domain Fusion with Weak Signals**:
   - Fused with 3 distant domains: language translation (data interpretation), gardening (controlled data pruning), music composition (sequence fine-tuning).
   - Generated 3 breakthrough concepts with Leap-Ahead Impact Statements:
     - Translation Fine-Tune AI: Fuses linguistic adaptation with AI for 12x efficient dataset fine-tuning in nuclear physics by 2028.
     - Gardening Prune Models: Applies plant care optimization to AI for 13x precise controlled fine-tuning on physics data.
     - Music Sequence AI: Integrates composition sequencing for 11x adaptive self-fine-tuning in stewardship models.
2. **Top Concept Selection**:
   - Most promising: Translation Fine-Tune AI.
   - One-line summary with quantified impact: Fuses language translation adaptation with AI to fine-tune nuclear physics models 12x more efficiently, reducing training time from weeks to days while maintaining accuracy.
   - Next-Best Action (blue sky version): Initiate unclassified proxy pilot using translation datasets at LANL within 3-6 months.
3. **Technical Correlation Explanation**:
   Cross-entropy loss $H(p,q) = -\sum p(x) \log q(x)$ minimizes divergence in translation fine-tuning (token probabilities) and nuclear dataset alignment (physics parameter distributions), allowing AI to transfer embedding spaces from linguistic corpora to quantum states. This correlation is non-obvious as translation handles semantic shifts in human languages, while nuclear fine-tuning manages probabilistic uncertainties in

particle interactions, yet both use gradient descent for alignment. Scientific validity from shared variational inference, enabling adapters to fine-tune across domains.

4. **SCARPETSE+G Score** (top concept only):
S=9, C=8, A=9, R=7, P=9, E=8, T=9, S=8, E=9, Global Equity=7, +G=9. Median=8.5 (≥8), no score <6. Flag: Regulation score reflects classified data handling in Stage 2.

5. **Vision Statement**:
Fine-tuning AI fused with language translation via cross-entropy optimization enables 12x efficient adaptation of nuclear physics models. This transfers interpretation techniques from open-source corpora to stewardship. The impact enhances stockpile accuracy in classified settings.

STAGE 2: UCPD v15.7 IMPLEMENTATION EXTRACTION

Context Bridge from Stage 1:
Current State: Manual fine-tuning on limited datasets using tools like TensorFlow on secure HPC, taking weeks.
Desired End State: 12x efficient fine-tuning via translation AI fusion.
Key Insight from Fusion: Cross-entropy bridges token to parameter alignment.

1. **Objective + BCI Hierarchy**:
Operational problem: Integrate open-source AI adapters into unclassified proxies for 12x fine-tuning efficiency in nuclear physics.
    - Classification breach: BCI = 10
    - Certification error: BCI = 10
    - Legacy system disruption: BCI = 8
    - Budget overrun >20%: BCI = 7
    - Cultural resistance: BCI = 5
    - Timeline slip >6 months: BCI = 6
2. **Driver Analysis**:
    - Accelerating: AI fine-tuning advancements (volatility: 0.39)
    - Stabilizing: NNSA Stockpile Stewardship protocols (volatility: 0.05)
    - Reversing: Dataset privacy challenges (volatility: 0.13)

3. **LSPE Output**:
   Configuration: Add adapter fine-tuning for cross-entropy to proxy models. Survives 2.5x amplification of budget and privacy forces. No CHL triggered.

4. **SMS Bridge**:

   - Install PEFT v0.13.0 on unclassified HPC. 2. Load GLUE via Datasets v3.0.1. 3. Fine-tune adapters on cross-entropy for 100 epochs. 4. Attach to surrogate inputs.

5. **Reversibility**:

   - Trigger conditions: Accuracy drop >4% or efficiency <10%.

   - Rollback steps: Remove adapters from config; revert base model; delete tuned weights.

   - Bounded cost: $450k + 110 hours.

6. **Golden Decision Brief**:
   Integrate PEFT v0.13.0 adapters using GLUE data into surrogates for unclassified proxy fine-tuning; justification is 20% efficiency gain; minimal adoption bridge is LoRA attachment; reversibility pathway is config removal with weight delete; validation checkpoints are 90/180-day parallels.

STAGE 3: SYNTHESIS FOR SUBMISSION

Title: Cross-Entropy AI for Nuclear Fine-Tuning

Executive Summary:
Language translation fused with AI via cross-entropy enables 12x efficient fine-tuning of nuclear physics models. Initial implementation integrates PEFT v0.13.0 into surrogates using GLUE for proxies. This achieves 20% gain in year 1, scaling to 12x by 2028.

Technical Approach:

- Current capability baseline: Weeks-long fine-tuning via TensorFlow on secure HPC.

- Proposed enhancement: Integrate PEFT v0.13.0 into surrogates using GLUE for proxies.

- Connection to breakthrough potential: Builds toward translation fusion for alignment insights, enabling non-obvious embedding transfer.

- Specific next steps with timeline: Month 1-3: Module development and integration; Month 4-6: Proxy validation; Month 7-12: Initial results analysis.

Risk Mitigation:

- Reversibility provisions: Remove adapters on >4% drop trigger, cost bounded at $450K + 110 hours.
- Parallel validation approach: Run base models concurrently for 180 days.
- Classification protection measures: Restrict to unclassified proxies; ITAR/EAR compliance via domestic development.

Partnerships & Resources:

- Lead organization: Los Alamos National Laboratory (LANL).
- Required collaborators: Five Eyes allies for data sharing; academic partners (e.g., MIT Nuclear Science).
- Estimated budget range: $1.5M-$2.5M for first year.
- Key personnel types needed: AI specialists, physicists, HPC administrators.

Success Metrics:

- 90-day milestone: Module integrated, initial proxy runs complete.
- 180-day milestone: 20% gain validated in analogs.
- 360-day go/no-go decision point: Accuracy equivalence to legacy with no errors; proceed to classified bridge if met.

Process improvement note: Expand fusion evidence with more arXiv specifics; adjust budgets based on tool-derived costs if available.

Click here to open the open source, AI assisted technical exploration and supplemental starter kit. *(Manual link:https://ucdpforge.com/pdfs/R2I/NNSA_07_IG_CrossEntropyAI_NukeLinguist.pdf)*