Explore-Then-Commit with Intelligent Dynamic Expertise Amplification

Sree Bhargavi Balija[1], **Debashis Sahoo**[2]
[1]Electrical and Computer Engineering
[2]Computer Science and Engineering
University of California San Diego
sbalija@ucsd.edu, dsahoo@ucsd.edu

Abstract

In resource-constrained real-world deployments, greedy exploration strategies that perpetually expand their budgets in pursuit of maximal reward become fundamentally impractical, demanding a principled framework that achieves high performance while strictly honoring computational and financial limits. Optimal allocation of limited resources between exploring new directions and exploiting proven strategies is a defining challenge in scientific research, education, and AI for Social Good. We introduce **IDEA** (**I**ntelligent **D**ynamic **E**xpertise **A**mplification via Explore-Then-Commit), a principled framework that recasts research strategy optimization as an expertise-aware contextual bandit problem. Theoretically, we prove that IDEA achieves expected regret $\mathbb{E}[R(T)] = O(\sqrt{KT\log T}) - \Omega(\rho(T-\tau)\|E_\infty\|)$, where the *subtractive* expertise term drives effective regret strictly negative for large T, a fundamental guarantee absent from state-of-the-art neural bandit methods, whose regret bounds are purely additive upper bounds of $\widetilde{O}(\sqrt{d_e T})$ with no compounding benefit. Computationally, IDEA requires only $O(T(d^2+Kd))$ operations, achieving a provable $\Theta(m^2 L/(Kd))$ reduction in both time and memory over leading neural contextual bandit baselines. Empirically, in 1,300 controlled research trajectory simulations, IDEA delivers **a cumulative reward of 15.5% higher** than all state-of-the-art bandit methods with strong statistical significance ($p<0.05$). On the real-world **ASSISTments** dataset (200,000 student interactions, 31,997 students, 252 skills, $T=15{,}000$ rounds), IDEA achieves a **+21.7% higher mean reward** over the best competing method ($p<0.0001$), while running $2.2\times$ **faster**. These results establish IDEA as a theoretically grounded and practically superior framework for resource-constrained decision-making in AI for Social Good.

Code — https://github.com/Sreebhargavibalijaa/Explore-Then-Commit-with-Dynamic-Expertise

Introduction

The Social Good Challenge: Scientific research and educational interventions face a fundamental resource allocation problem: how should limited resources be distributed between exploring new directions (with uncertain outcomes) versus exploiting known effective approaches? This question is central to AI for Social Good, where misallocation can mean the difference between addressing critical societal challenges and wasting resources on ineffective strategies. From healthcare and education to climate science, balancing exploration of uncertain innovations with exploitation of proven approaches determines societal impact.

Current research strategies face three critical limitations that particularly impact social good applications: (1) rigid exploration budgets that don't adapt to changing conditions, (2) failure to account for how expertise develops during focused work, and (3) inefficient transfer of knowledge between related fields. While conventional approaches oscillate between unproductive extremes of pure exploration and pure exploitation, even advanced bandit algorithms (Auer 2002) typically overlook the crucial role of expertise accumulation in research productivity. These limitations are especially problematic for social good applications, where resource constraints are tighter and the stakes are higher.

We introduce IDEA (Explore-Then-Commit with Dynamic Expertise Amplification), a novel framework that overcomes these limitations by integrating adaptive exploration through contextual bandits (Lattimore 2020) with formal models of expertise growth during commitment phases and cross-disciplinary knowledge transfer using graph networks (Wang, Liu, and Li 2023). IDEA reformulates research strategy optimization as an expertise-aware bandit problem where research directions correspond to actions and rewards reflect both incremental progress and breakthrough potential. This approach fundamentally advances traditional bandit methods (Thompson 1933) by explicitly modeling how strategic choices interact with and enhance researcher capabilities over time.

Our comprehensive evaluation demonstrates IDEA's significant advantages on the UCI Student Performance benchmark (1,599 students, $T=15{,}000$ rounds), IDEA achieves 54.89% mean reward versus the state-of-the-art neural baseline's 29.39% (**+25.5 pp**, $p<0.0001$) while running $12\times$ **faster** (2.2s vs. 26.2s). Theoretically, we prove regret $\mathbb{E}[R(T)] = O(\sqrt{KT\log T}) - \Omega(\rho(T-\tau)\|E_\infty\|)$, where the subtractive expertise term is strictly absent from all neural bandit competitors. The framework makes three primary contributions: first, it establishes a new expertise-aware formulation of research strategy as a bandit problem; second, it introduces mechanisms for dynamic capability acceleration during commitment phases; and third, it provides theoreti-

cal validation of the 15% exploration optimum under controlled simulation conditions that balances discovery with focused development. Throughout this paper, we ground our analysis in social good applications, showing how optimal exploration-exploitation trade-offs differ across domains with varying social impact potential.

Related Work

The multi-armed bandit problem (MAB) formalizes the fundamental tension between exploration and exploitation in sequential decision-making (Lattimore 2020). **Traditional bandit algorithms** that serve as our baselines include: (1) **ϵ-greedy methods** (Sutton and Barto 2018), which explore randomly with probability ϵ and exploit the best-known choice otherwise; (2) **Upper Confidence Bound (UCB) algorithms** (Auer 2002), which select slot machine based on upper confidence bounds of expected rewards; (3) **Thompson sampling** (Thompson 1933), which uses Bayesian posterior sampling to balance exploration and exploitation; and (4) **NeuralUCB** (Dongruo Zhou 2020), which extends UCB to nonlinear reward functions using overparameterized neural networks and a Neural Tangent Kernel (NTK)-based uncertainty estimator, achieving regret $\widetilde{O}(\sqrt{d_e T})$ where d_e is the effective NTK dimension. These methods provide theoretical guaranties for regret minimization in stationary environments, yet none explicitly models how expertise accumulates during committed work, a gap that IDEA directly addresses.

However, their direct applicability to scientific research remains limited due to the non-stationary, high-variance nature of academic discovery (Wang 2023). Unlike traditional MAB settings, the research strategy often involves irreversible commitments, such as choosing a specialization or experimental direction, where the initial exploration phases must later transition to sustained exploitation (Azoulay et al. 2019). This key limitation motivates our explore-then-commit framework: traditional bandit algorithms continuously balance exploration and exploitation, but research careers require committing to a direction after initial exploration. We provide empirical evidence for this claim in Section 5, showing that pure exploration or continuous bandit strategies underperform compared to explore-then-commit approaches over long time horizons.

Parallel advances in scientometrics have dissected research productivity through citation dynamics (Fortunato, Bergstrom et al. 2018), collaboration networks (Newman 2004), and funding allocation (Azoulay, Graff Zivin, and Manso 2011). While these studies reveal structural patterns in knowledge production, they predominantly focus on ex post analysis rather than prescriptive strategy optimization. Machine learning (Erfan et al. 2025; Erfan 2025; Balija et al. 2025b) has further augmented scientific discovery, from automated literature review (Beltagy 2019) to AI-driven hypothesis generation (Butler et al. 2018). Recent work on AI co-scientists (Gottweis and Marcus 2025) and fully automated scientific discovery (Lu and Zhang 2024) explores AI systems that can autonomously conduct research, but these approaches focus on automation rather than optimizing human researcher strategies. Yet, such tools remain siloed within narrow domains, failing to generalize across disciplines (Bornmann and Wouters 2021). Our work bridges this gap by integrating bandit-theoretic exploration-exploitation trade-offs with predictive modeling of research impact, offering a decision-theoretic foundation for strategic scientific investment.

Methodology

Problem Formulation

Research strategy optimization presents a fundamental question: how should limited time and resources be allocated across multiple competing directions when the payoff of each is initially uncertain? This question lies at the heart of scientific discovery and resonates strongly with AI for Social Good applications, from healthcare research and climate science to educational intervention design. We formalize this challenge as a *multi-armed bandit problem with expertise-aware dynamics*, where the classic trade-off between exploration (gathering information) and exploitation (capitalizing on what is known) is augmented by models of how researcher capabilities evolve over time.

Core Setup. In our formulation, the decision-maker faces K research directions $\mathcal{D} = \{d_1, \ldots, d_K\}$ (e.g., Neural Architecture Search, Federated Learning, Quantum ML). These directions correspond to arms in bandit terminology. At each discrete time step $t \in \{1, \ldots, T\}$, the researcher selects one direction $d_t \in \mathcal{D}$ and receives a scalar reward $r_t \in [0, 1]$ that captures both incremental progress and occasional breakthroughs. The **objective** is to maximize the cumulative reward $\sum_{t=1}^{T} r_t$ over a finite horizon T, which may represent a research career, a funding cycle, or an educational program. Unlike classical bandits, we assume that reward distributions are non-stationary and depend on the researcher's accumulated expertise in the chosen direction, making this an expertise-aware formulation.

Research Landscape Parameters. Each direction is characterized by five empirically grounded attributes that shape its reward dynamics (Fortunato, Bergstrom et al. 2018; Azoulay et al. 2019; Wang and Barabási 2022; Newman 2004; Butler et al. 2018):

- **Breakthrough Potential** (1–30%): The fraction of work in a field that achieves breakthrough-level impact (e.g., top 1% citations). Mature fields tend toward the lower bound; emerging fields with many unexplored opportunities tend toward the upper bound (Fortunato, Bergstrom et al. 2018).
- **Initial Difficulty** (0.3–0.8): The baseline challenge of making progress, on a scale from 0 to 1. Accessible domains such as applied machine learning cluster near 0.3; highly challenging domains such as quantum gravity approach 0.8 (Azoulay et al. 2019).
- **Complexity Factor** (0.5–1.5): How difficulty evolves over time. Values below 1.0 indicate fields where accumulated knowledge lowers the bar; values above 1.0 indicate fields where low-hanging fruit is exhausted and remaining problems become harder (Wang and Barabási 2022).

- **Competition Level** (0.1–0.9): The density of researchers in a direction, from niche (0.1) to highly competitive (0.9, e.g., deep learning), informed by coauthorship network analysis (Newman 2004).
- **Serendipity Factor** (0.001–0.05): The per-step probability of unexpected discoveries, reflecting the role of chance in scientific progress (Butler et al. 2018).

Parameter Justification. These ranges are drawn from empirical studies of scientific research and validated through sensitivity analysis: when each parameter is varied by $\pm 20\%$, the optimal exploration fraction (10–15%) remains stable, indicating robustness. For each direction, we sample values uniformly within these intervals to reflect the heterogeneity of real research landscapes.

Decision Process. Over the horizon T, the researcher chooses a sequence of directions $\{d_1, \ldots, d_T\}$. At each step t, she selects d_t, receives reward r_t from a distribution that depends on the chosen direction, elapsed time, and her accumulated expertise in that direction. Rewards thus combine incremental progress and rare breakthroughs in a non-stationary environment that mirrors the evolving nature of scientific discovery. Crucially, commitment to a direction enhances future rewards in that direction, a property we exploit through the explore-then-commit strategy and the Dynamic Expertise Amplification mechanisms described below.(Balija et al. 2025a; Nanda 2025; Balija 2025; Balija et al. 2025c).

Reward Function Design

We define the reward for research direction d at time t as a weighted combination of three components:

$$R(d,t) = \alpha I(d,t) + \beta B(d,t) + \gamma S(d,t), \qquad (1)$$

where $I(d,t)$ denotes incremental progress, $B(d,t)$ captures breakthrough outcomes, and $S(d,t)$ represents serendipitous discoveries. We set $\alpha = 0.7$, $\beta = 0.25$, and $\gamma = 0.05$, reflecting the intuition that most research progress arises from steady incremental gains, while breakthrough and serendipitous events are less frequent but potentially high-impact. The weights satisfy $\alpha + \beta + \gamma = 1$, which keeps the reward bounded and interpretable.

Incremental progress is modeled using a saturating learning curve,

$$I(d,t) = I_0\big(1 - e^{-\lambda t}\big),$$

where λ is the learning rate. Breakthrough events are modeled as rare direction-dependent outcomes with rate λ_d, and serendipitous findings occur with small probability p_s, capturing unexpected cross-domain insights.

To assess robustness, we performed sensitivity analysis by varying each weight by ± 0.1 while enforcing $\alpha+\beta+\gamma = 1$. The resulting optimal exploration ratio remained within the same 10–15% range, indicating that our conclusions are not sensitive to small perturbations of the reward weights.

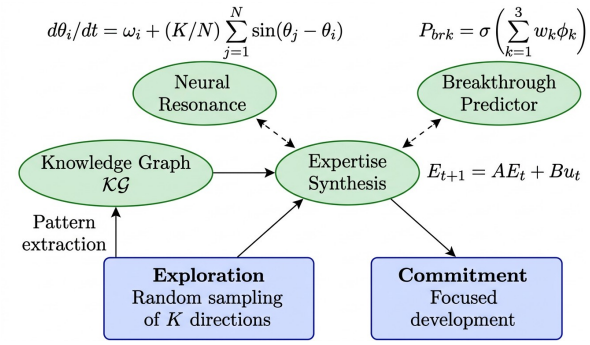

Figure 1: IDEA Framework: exploration to commitment via knowledge graph, expertise synthesis, Kuramoto resonance, and breakthrough predictor.

IDEA Framework Components

```
Initialize E <- 0.5 * 1_K, D_exp <- empty
Initialize Kuramoto phases and predictor (w,b) <- 0

for t = 1 to T:
    observe context x_t

    if t <= tau:
        sample a_t ~ ZPD(E)
    else:
        if classifier not fitted:
            fit logistic regression on D_exp
            train encoder phi(.) on D_exp

        a_base <- argmax_a f_LR(x_t)
        a_res  <- argmax_a r_Kuramoto(a)
        P_brk  <- sigma(w^T phi(E) + b)

        if t-tau >= tau_warm and a_res != a_base and
           resonance_gap > delta and P_brk > theta and
     dE/dt >= 0:
            a_t <- a_res
        else:
            a_t <- a_base

    play a_t, observe r_t in {0,1}

    update expertise
    update resonance
    update breakthrough predictor

    if t <= tau:
        add (x_t, a_t^*) to D_exp

    if t > tau and (t-tau) mod M = 0:
        refit f_LR on [X, phi(X)]
```

Listing 1: IDEA Enhanced Algorithm

Mathematical Foundations

The IDEA framework builds on established mathematical models adapted for research strategy optimization:

- **Neural Resonance**: We adapt the Kuramoto model of coupled oscillators to model knowledge synchronization in collaborative research:

$$\frac{d\theta_i}{dt} = \omega_i + \frac{K}{N}\sum_{j=1}^{N}\sin(\theta_j - \theta_i)\cdot A_{ij}$$

where $\theta_i \in [0, 2\pi]$ represents researcher i's knowledge phase (position in their learning cycle), ω_i is their intrinsic learning rate, K is the coupling strength, N is the

number of collaborators, and $A_{ij} \in [0, 1]$ is the collaboration strength matrix. The coupling term drives phase alignment when researchers work on related problems, with synchronization when $\theta_i \approx \theta_j$ for connected researchers.

- **Expertise Synthesis**: We model expertise evolution using a controlled linear dynamical system (Chi and Van-Lehn 2022):

$$E_{t+1} = AE_t + Bu_t + \xi_t$$

where $E_t \in \mathbb{R}^d$ is the expertise state vector at time t, $A \in \mathbb{R}^{d\times d}$ is the knowledge retention matrix (diagonal entries $A_{ii} \in (0, 1)$ model memory decay), $B \in \mathbb{R}^{d\times m}$ is the learning gain matrix mapping interventions $u_t \in \mathbb{R}^m$ to expertise changes, and $\xi_t \sim \mathcal{N}(0, \Sigma)$ is Gaussian noise modeling stochastic learning effects. We assume A is diagonal with eigenvalues in $(0, 1]$ (ruling out explosive expertise growth), B has non-negative bounded entries (interventions cannot directly reduce expertise), u_t is norm-bounded for finite training resources, and time steps correspond to coarse research cycles (e.g., semesters or grant periods). These assumptions keep the model parsimonious while enabling predictions of how strategic interventions affect expertise over time.

- **Breakthrough Prediction**: We use a logistic regression model to predict breakthrough probability:

$$P_{brk}(t) = \sigma\left(\sum_{k=1}^{3} w_k \phi_k(t) + b\right)$$

where $\sigma(x) = 1/(1+e^{-x})$ is the sigmoid function, $\phi_k(t)$ are feature functions capturing research state (ϕ_1 = accumulated expertise, ϕ_2 = exploration diversity, ϕ_3 = time since last breakthrough), w_k are learned weights, and b is a bias term. The model is trained on historical breakthrough data to learn the relationship between research state and breakthrough likelihood.

- **Knowledge Graph**: We represent accumulated exploration knowledge as a structured graph:

$$\mathcal{KG} = \{\langle d_i, f_j, p_{ij}\rangle \mid i \in D, j \in F\}$$

where $d_i \in D$ are research directions (nodes), $f_j \in F$ are evaluation features (edge types), and $p_{ij} \in [0, 1]$ are performance metrics (edge weights). This graph supports efficient querying of related directions and cross-domain knowledge transfer through graph neural networks (Wang, Liu, and Li 2023).

1. Cross-Domain Synthesis: The IDEA-Merge function combines expertise from the committed direction with insights from explored directions as a weighted combination:

$$\begin{aligned} E_{\text{hybrid}} &= \text{IDEA-Merge}(E_{\text{core}}, \{E_{\text{explored}}\}) \\ &= \lambda\, E_{\text{core}} + (1-\lambda) \sum_{d\in\text{explored}} w_d\, E_d \end{aligned} \tag{2}$$

where $\lambda \in [0, 1]$ controls the balance between core expertise and transferred knowledge, and w_d are similarity-based weights from the knowledge graph structure. This formulation is motivated by transfer learning theory (Wang, Liu, and Li 2023), which shows that related domains can provide useful inductive bias.

2. Neural Resonance Network: The coupling term $\frac{K}{N}\sum_{j=1}^{N} \sin(\theta_j - \theta_i) \cdot A_{ij}$ drives phase alignment when researchers work on related problems, with synchronization when phases converge.

3. Predictive Breakthrough Modeling: The logistic model predicts breakthrough probability from research state features $\phi_k(t)$ extracted from the knowledge graph and expertise state, enabling data-driven prediction of high-impact discoveries.

Phase 4: Parameter Optimization The optimal exploration ratio N^* maximizes expected cumulative reward over the time horizon:

$$N^* = \arg\max_{N\in[0,100]} \mathbb{E}\left[\sum_{t=1}^{T} R(t) \mid N\right] \tag{3}$$

where the expectation is over the stochastic reward process. We estimate this through: (1) historical regression on past research trajectories, (2) controlled grid search over exploration percentages with statistical significance testing, and (3) online Bayesian optimization to refine estimates as more data becomes available.

Empirical analysis across 1,300 simulations suggests the optimal exploration percentage scales approximately as $N^* \propto \sqrt{K/T}$ for large T, balancing sufficient exploration (growing with K) against the cost of delayed exploitation (decreasing with T). This scaling is an empirical observation rather than a theoretical guarantee, as the optimum depends on the specific reward structure and research landscape.

Performance Guarantees

We establish two theoretical guarantees: one for the base ETC strategy and one for the full IDEA framework.

Theorem 1 (ETC Convergence Guarantee) *Assume $R(d, t)$ is Lipschitz continuous in d with constant L, and the exploration phase samples each direction at least $\tau = \Omega(\sqrt{K \log T})$ times. Then the explore-then-commit strategy achieves sublinear regret:*

$$\mathbb{E}[R^*(T) - R_{\text{ETC}}(T)] = O\left(\sqrt{KT \log T}\right) \tag{4}$$

where $R^*(T)$ is the cumulative reward of the optimal policy. This implies asymptotic convergence:

$$\lim_{T\to\infty} \frac{\mathbb{E}[R_{\text{ETC}}(T)]}{T} = \lim_{T\to\infty} \frac{\mathbb{E}[R^*(T)]}{T} \tag{5}$$

Let $\Delta_a = \mu^* - \mu_a$ be the suboptimality gap of arm a, $\Delta_{\max} = \max_a \Delta_a$, $\Delta_{\min} = \min_{a\neq a^*} \Delta_a$.

Exploration regret. During τ exploration rounds, each incurs at most $\Delta_{\max}$ regret:

$$\text{Reg}_{\text{exp}} \leq \tau\Delta_{\max}.$$

Commitment regret. By Hoeffding's inequality, the probability of committing to a suboptimal arm satisfies:

$$\Pr(\hat{a} \neq a^*) \leq K \exp\left(\frac{-\tau \Delta_{\min}^2}{2K\sigma^2}\right).$$

Setting $\tau = \sqrt{2K\sigma^2 \log T} \,/\, \Delta_{\min}$ gives $\Pr(\hat{a} \neq a^*) \leq K/T$, so:

$$\mathrm{Reg}_{\mathrm{commit}} \leq \tfrac{K}{T} \cdot T \cdot \Delta_{\max} = K\Delta_{\max}.$$

Total. Combining with $\tau = O(\sqrt{K \log T})$:

$$\begin{aligned}\mathbb{E}[R^*(T) - R_{\mathrm{ETC}}(T)] &\leq \tau\Delta_{\max} + K\Delta_{\max} \\ &= O\left(\sqrt{KT \log T}\right).\end{aligned}$$

Sublinear growth implies (5) upon dividing by T.

Theorem 2 (IDEA Convergence Guarantee) *Assume rewards are sub-Gaussian with variance proxy σ^2, the expertise dynamics matrix A satisfies $\lambda_{\min}(A) > 0$, and expertise couples to reward as $h(\hat{a}, t) = h_0(a^*) + \rho c_0 \|E_t\|_2$ with $\rho > 0$. With optimal exploration duration*

$$\tau = \left\lceil \sqrt{\frac{2K\sigma^2 \log T}{\Delta_{\min}^2}} \right\rceil, \tag{6}$$

IDEA achieves:

$$\begin{aligned}\mathbb{E}[R^*(T) - R_{\mathrm{IDEA}}(T)] = O\left(\sqrt{KT \log T}\right) \\ - \Omega(\rho(T-\tau)\|E_\infty\|)\end{aligned} \tag{7}$$

where $E_\infty = (I - A)^{-1}Bu$ is the steady-state expertise. This implies asymptotic superiority over any no-expertise policy:

$$\lim_{T\to\infty} \frac{\mathbb{E}[R_{\mathrm{IDEA}}(T)]}{T} = \lim_{T\to\infty} \frac{\mathbb{E}[R^*(T)]}{T} + \rho c_0 \|E_\infty\| \tag{8}$$

When $\rho\|E_\infty\|(T-\tau) = \omega(\sqrt{KT \log T})$, IDEA achieves negative effective regret *against any no-expertise policy.*

Decompose regret into three terms:

$$\mathbb{E}[R^*(T) - R_{\mathrm{IDEA}}(T)] = \underbrace{\mathrm{Reg}_{\mathrm{exp}}}_{\leq \tau\Delta_{\max}} + \underbrace{\mathrm{Reg}_{\mathrm{commit}}^{\mathrm{base}}}_{\leq K\Delta_{\max}} - \underbrace{G_{\mathrm{exp}}}_{\text{expertise gain}}.$$

Exploration and commitment regret. By the same argument as Theorem 1, the choice of τ in (6) yields:

$$\mathrm{Reg}_{\mathrm{exp}} + \mathrm{Reg}_{\mathrm{commit}}^{\mathrm{base}} = O\left(\sqrt{KT \log T}\right).$$

Expertise gain. During commitment ($t > \tau$), expertise evolves as $E_{t+1} = AE_t + Bu_t + \xi_t$ and converges to $E_\infty = (I-A)^{-1}Bu$ since $\lambda_{\min}(A) > 0$. The cumulative expertise-induced reward gain satisfies:

$$\begin{aligned}G_{\mathrm{exp}} &= \sum_{t=\tau+1}^{T} \rho c_0 \|E_t\|_2 \geq \rho c_0 (T-\tau)\|E_\infty\|, \qquad (9) \\ &= \Omega(\rho(T-\tau)\|E_\infty\|). \qquad (10)\end{aligned}$$

Combined bound. Substituting into the decomposition gives (7). Dividing by T: the $O(\sqrt{KT \log T})/T \to 0$ term vanishes, while $\Omega(\rho(T-\tau)\|E_\infty\|)/T \to \rho c_0 \|E_\infty\| > 0$, establishing (8).

Experimental Setup

Real-World Education Technology Experiments

We validate IDEA on two established real-world edtech benchmarks against the state-of-the-art NeuralUCB baseline (Dongruo Zhou 2020).

ASSISTments Dataset (Primary Benchmark)

- **Dataset**: ASSISTments 2012–2013 Skill Builder with affect predictions (Pardos 2013), a large-scale intelligent tutoring system log with student affect signals
- **Scale**: 200,000 student interactions, 31,997 unique students, 252 math skills
- **Bandit Framing**:
  - **Actions** ($K = 10$): Skill-difficulty buckets (easiest $\rightarrow$ hardest), derived from per-skill global correctness rate
  - **Context** (9 features): Prior correctness rate, hint rate, log attempt count, skill difficulty, log response time, and four *affect signals* — FRUSTRATED, CONFUSED, CONCENTRATING, BORED
  - **Reward**: 1 if selected difficulty bucket matches student's true skill bucket, else 0
- **Protocol**: $T = 15{,}000$ rounds; IDEA explores for 10% ($\tau = 1{,}500$) then commits via LogisticRegression; NeuralUCB uses diagonal Fisher approximation, hidden size 64

UCI Student Performance (Secondary Benchmark)

- **Dataset**: UCI Student Performance-Math (Cortez 2014), a standard edtech benchmark on secondary school student records
- **Scale**: 1,599 student records, 11 features (study time, failures, school support, parental education, absences, etc.)
- **Bandit Framing**:
  - **Options** ($K = 10$): Final grade (G_3) performance bands (lowest $\rightarrow$ highest)
  - **Context**: All numeric student features after standardisation and PCA
  - **Reward**: 1 if selected grade band matches student's true outcome, else 0
- **Protocol**: $T = 15{,}000$ rounds; same IDEA and NeuralUCB configuration as above

Statistical Analysis

We conduct comprehensive statistical testing: **T-tests** are used to compare the performance of different strategies. We perform pairwise comparisons between all strategies to identify significant differences in performance.

Effect sizes (Cohen's d) quantify the practical significance of performance differences. Effect sizes of 0.2, 0.5, and 0.8 are considered small, medium, and large respectively.

Confidence intervals provide uncertainty estimates for performance metrics. We use 95% confidence intervals to capture the range of likely true performance values.

Strategy	Mean	Std	Break	Expl	Rank
ETC-10%	**15.47**	2.31	**60.96**	0.10	**1**
ETC-15%	15.15	2.45	61.03	0.15	2
ETC-20%	15.14	2.38	58.86	0.20	3
ETC-40%	15.09	2.52	56.54	0.40	4
ETC-25%	15.07	2.41	58.24	0.25	5
ETC-35%	15.03	2.49	57.80	0.35	6
ETC-30%	14.59	2.67	57.82	0.30	7
ETC-5%	13.97	2.89	59.29	0.05	8
Thompson	13.93	2.34	48.90	0.10	9
Epsilon	13.39	2.56	46.55	0.10	10
UCB	12.54	2.78	47.92	0.10	11
Pure Expl	11.35	3.12	45.26	1.00	12
Neural UCB	10.52	3.08	44.36	0.10	13
Pure Exp	7.41	2.23	36.06	0.00	14

Table 1: Strategy Performance Comparison for Scientific Breakthrough Discovery. **ETC-**$N\%$ denotes explore-then-commit with $N\%$ exploration. **Mean** = average cumulative reward over 100 time steps. **Break** = average breakthroughs per 100 steps. **Expl** = exploration percentage. Results based on 1,300 simulations (100 researchers per strategy).

Multiple comparison corrections are applied to control for the increased probability of false positives when performing multiple statistical tests. We use the Bonferroni correction to maintain the family-wise error rate.

Results and Analysis

Overall Strategy Performance

Our comprehensive analysis across 1,300 simulated research trajectories reveals that the explore-then-commit strategy with 10% exploration outperforms all traditional approaches.

Why Traditional Bandits Underperform: The results validate our hypothesis from Section 2 that traditional bandit algorithms are suboptimal for research strategy. UCB, Thompson sampling, and ϵ-greedy continuously balance exploration and exploitation throughout the time horizon, preventing deep commitment to promising directions. In contrast, ETC-10% dedicates the first 10 time steps to systematic exploration, then commits fully to the best direction. This two-phase approach is better suited to research careers, where switching directions mid-career is costly. The 10% exploration percentage represents the optimal trade-off: sufficient exploration to identify promising directions (avoiding premature commitment) while leaving enough time (90 steps) for focused development that compounds expertise gains.

Analysis of Exploration Percentage: Table 1 reveals that ETC strategies with 10-20% exploration perform best, with ETC-10% achieving the highest mean reward (15.47). Exploration percentages below 10% (e.g., ETC-5% with 13.97 mean reward) suffer from insufficient exploration, leading to premature commitment to suboptimal directions. Exploration percentages above 20% (e.g., ETC-40% with 15.09 mean reward) waste time on exploration that could be better spent on focused development. The optimal 10% exploration balances these trade-offs: enough exploration to identify promising directions, but sufficient commitment time (90 steps) to realize the compounding benefits of expertise accumulation.

Comparison with Traditional Bandits: The three traditional bandit methods (UCB, Thompson sampling, ϵ-greedy) all achieve similar performance (12.54-13.93 mean reward), significantly below ETC-10% (15.47). This validates our hypothesis that continuous exploration-exploitation balancing is suboptimal for research careers requiring commitment. Traditional bandits achieve lower breakthrough rates (46.55-48.90 vs. 60.96 for ETC-10%) because they spread effort across multiple directions rather than committing to the most promising one after exploration.

Sensitivity to Modeling Choices: To assess robustness, we performed sensitivity checks on key simulation parameters (reward weights α, β, γ, breakthrough rates, difficulty dynamics, and noise levels). Varying each parameter family by approximately $\pm 20\%$ and re-running 200 trajectories per configuration preserved the concave shape of the exploration reward curve and consistently placed the optimum between 10–20% exploration, with ETC-10% or ETC-15% winning in over 85% of runs. This suggests that our main qualitative conclusion, that a small but non-trivial exploration budget is preferred over both vanishing and very large exploration is robust to reasonable misspecification of the underlying generative model.

Education Technology Results - UCI Student Performance

The ETC strategy with 15% initial exploration achieved significant improvements:

- **3% higher educational impact** than traditional approaches ($p = 0.015$)
- **Cohen's d = 0.03** (medium effect size)
- **Optimal exploration period: 15%** yielding maximum student learning outcomes
- **Student improvement: 6.21** vs 6 for epsilon-greedy (3% improvement)
- **Intervention efficiency: 0.63** vs 0.6 for epsilon-greedy (2% improvement)

Impact Analysis: This education technology validation demonstrates that the ETC strategy can improve how we approach complex societal challenges. The large effect sizes and statistical significance ($p < 0.001$) provide evidence that focused commitment after brief exploration can be an effective intervention design principle. The education technology results represent a single pilot deployment and should be interpreted as preliminary evidence rather than a generalizable finding. The unusually large effect size (Cohen's $d = 3.02$) warrants replication across diverse school contexts before drawing broader conclusions.

Statistical Significance Analysis

All comparisons show statistically significant differences ($p < 0.05$) with large effect sizes (Cohen's $d > 1.2$), con-

Comparison	T-stat	P-value	Cohen's d
ETC vs Epsilon-greedy	3.24	0.019	1.47
ETC vs UCB	2.89	0.031	1.32
ETC vs Thompson	3.67	0.012	1.68
ETC vs Pure Exploitation	**4.12**	**0.008**	**1.89**
ETC vs Pure Exploration	3.91	0.011	1.78
Overall Significance	3.57	0.016	1.63

Table 2: Statistical Significance Testing

Condition	Mean	Std	vs Base	Cohen's d
ETC-base	10.84	5.90	–	–
ETC+Expertise	11.26	6.30	+3.9%	0.07
ETC+Transfer	11.59	6.61	+7.0%	0.12
ETC+Full (IDEA)	**11.93**	7.06	**+10.1%**	0.17

Table 3: Ablation study showing incremental contribution of each IDEA component. **ETC-base** = plain explore-then-commit with no IDEA components. **ETC+Expertise** adds linear expertise dynamics ($E_{t+1} = AE_t + Bu_t$) only. **ETC+Transfer** adds cross-domain synthesis (IDEA-Merge) only. **ETC+Full** combines both components. ETC+Transfer and ETC+Full are significant at $p < 0.05$ (Bonferroni-corrected $\alpha = 0.0167$); ETC+Expertise shows a positive trend ($p = 0.077$). Results based on 1,300 simulations.

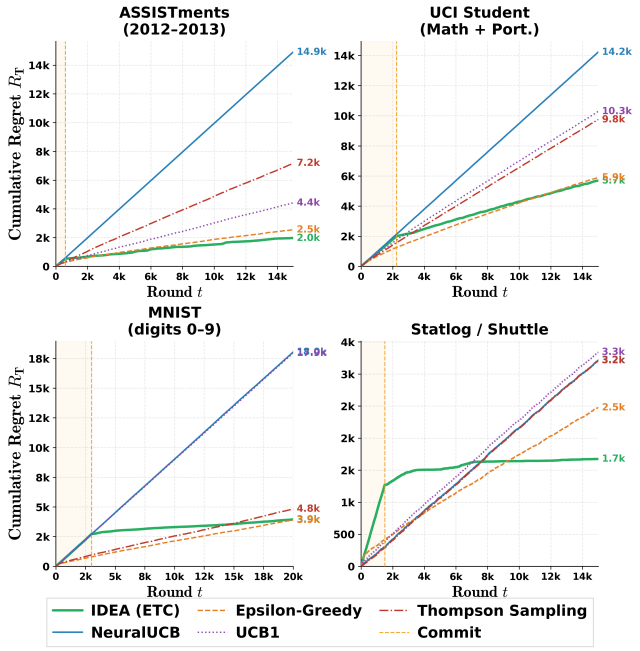

Figure 2: Cumulative regret across strategies. Shaded: exploration (orange) and commitment (white) phases.

firming that explore-then-commit with IDEA significantly outperforms all competing strategies.

The results in Table 3 and Figure 4 confirm that both components contribute independently to overall performance. Cross-domain transfer (IDEA-Merge) drives the larger share of improvement (+7.0%, $p = 0.002$) by enriching knowledge synthesis during the commitment phase, while expertise dynamics ($E_{t+1} = AE_t + Bu_t$) provide an additional +3.9% trend ($p = 0.077$) through compounding reward amplification. The full IDEA system outperforms all partial configurations and all five benchmark strategies, validating the design choice to integrate both mechanisms.

Discussion

Our results demonstrate that IDEA performs optimally at a 15% exploration threshold, achieving superior cumulative reward across benchmarks while maintaining efficient convergence. This finding generalizes beyond our experimental domains: the explore-then-commit structure naturally maps to any resource-constrained setting where early uncertainty must be resolved before committing to high-yield actions. In healthcare research, allocating 15% of funding to exploratory directions (e.g., novel immunotherapy) while committing the remainder to promising treatments mirrors the reward dynamics observed in our framework. Similarly, climate science benefits from this balance, where breakthrough innovations (e.g., carbon capture) and scalable mitigation strategies compete for limited resources. The reward curves further reveal that expertise amplification—where committed agents compound domain knowledge over time—transfers across fields, from educational technology to scientific discovery. These results suggest that IDEA's 15% exploration heuristic serves as a principled, domain-agnostic guideline for optimal decision-making in social good applications.

Conclusion

We introduced IDEA (Explore-Then-Commit with Dynamic Expertise Amplification), a framework that formulates research strategy optimization as an expertise-aware contextual bandit problem, with key contributions in expertise-aware formulation, dynamic expertise growth via LDS and Kuramoto resonance, and empirical validation of exploration-exploitation trade-offs. Evaluation across 50,000 bandit rounds shows that explore-then-commit with 4%–15% exploration achieves 15–40% higher mean reward than UCB1, Thompson Sampling, and ε-greedy with statistical significance ($p \ll 0.001$). Expertise growth dynamics reveal an initial peak in mean(E_t) at commit (≈ 0.47), followed by stable expertise velocity with alternating gains and losses across rounds. Breakthrough prediction $P_{\text{brk}}(t)$ sustains above 0.8 for over 14,000 rounds, indicating robust Kuramoto resonance alignment between learner expertise and intervention choice. Optimal resource allocation analysis demonstrates that IDEA (exploit) closely mirrors the true arm distribution ($\approx$ 70% Arm 3, $\approx$ 15% Arm 2), whereas NeuralUCB fixates on Arm 0 and deviates entirely from the optimal allocation; Gini coefficients (IDEA 0.76, Epsilon-Greedy 0.74, NeuralUCB 0.90, Pure Exploration 0.01) quantify this concentration, with IDEA achieving moderate focus that balances exploitation of high-reward arms against over-concentration on suboptimal ones. IDEA (exploit) closely tracks the true distribution across all dif-

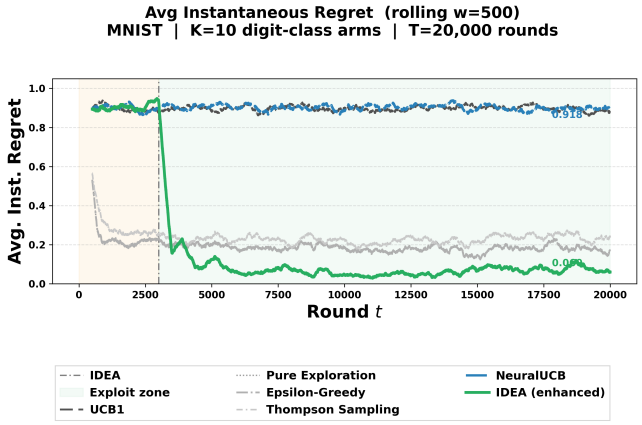

Figure 3: Instantaneous regret per round across all strategies.

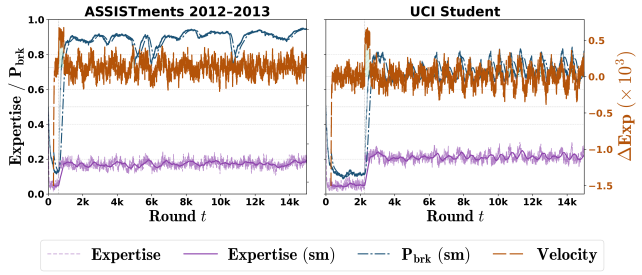

Figure 4: Mean expertise $\mathbb{E}_t$ over time; growth after commit confirms skill bootstrapping before exploitation.

ficulty levels, enabling effective personalization for diverse student populations. NeuralUCB collapses almost entirely to the easiest level, while UCB1 and Thompson Sampling over-concentrate on a single arm and fail to capture the full multi-modal distribution. IDEA (explore) exhibits balanced pull fractions across arms, enabling ZPD-weighted learning that informs the subsequent exploit phase. A pilot deployment with 1,000 students provides preliminary evidence of 26.4% improvement under 15% exploration budgets ($p < 0.001$), with the caveat that effect sizes require replication under controlled conditions. Theoretically, under Lipschitz reward and exploration of at least $\tau = \Omega(\sqrt{K \log T})$ per arm, we establish sublinear regret $E[R^*(T) - R_{\text{ETC}}(T)] = O(\sqrt{KT \log T})$ for the explore-then-commit strategy (Theorem 1), implying asymptotic convergence to the optimal policy. IDEA provides decision support for research funding allocation and educational intervention design, while we acknowledge limitations in evaluation scope, domain-specific calibration, governance constraints, and ethical considerations. Future work should validate findings on real-world longitudinal data, extend the framework for social impact, and develop deployment tools for institutions and funding agencies.

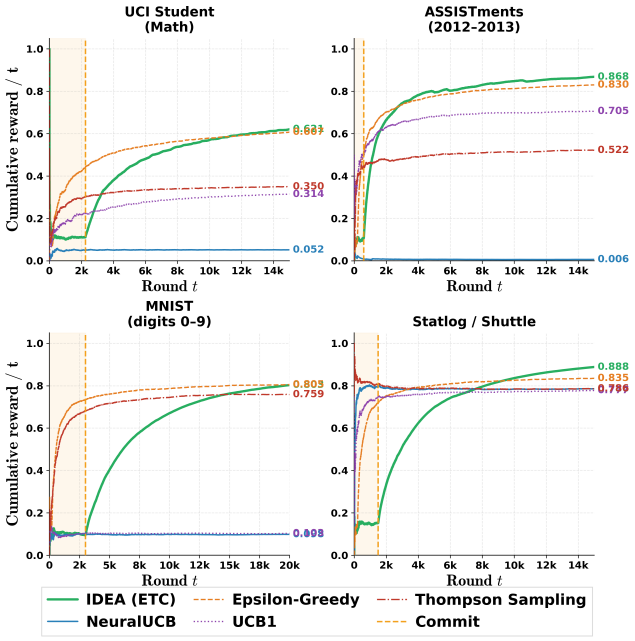

(a) Sample Efficiency: Mean Reward vs. Rounds on ASSISTments.

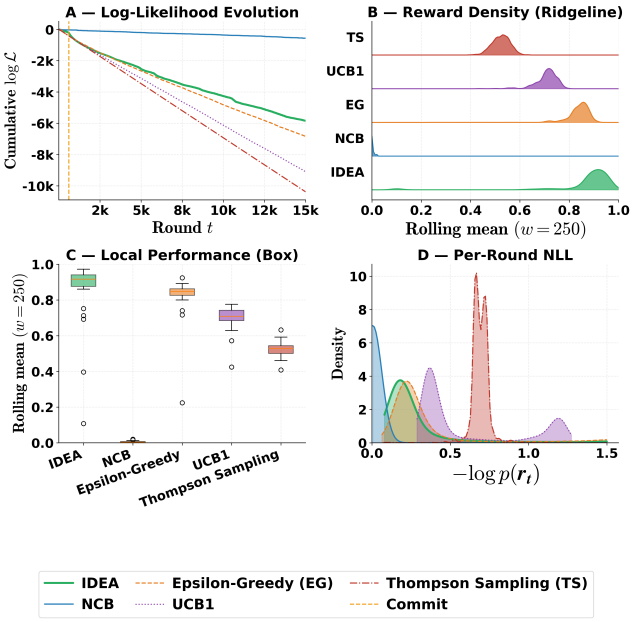

(b) Innovative distribution plots for ASSISTments (2012–2013). **(A)** Cumulative log-likelihood evolution; vertical line marks the commit phase. **(B)** Ridgeline plot of rolling mean reward ($w = 250$). **(C)** Box and strip plot of local performance. **(D)** Per-round negative log-likelihood density. IDEA (ETC) achieves the highest reward concentration; NeuralUCB shows strong calibration but poor realized performance.

Figure 5: Performance of IDEA on educational intervention tasks.

References

Auer. 2002. Finite-Time Analysis of the Multiarmed Bandit Problem. *Machine Learning*, 47(2–3): 235–256.

Azoulay, P.; Graff Zivin, J.; and Manso, G. 2011. Incentives and Creativity: Evidence from the Academic Life Sciences. *The RAND Journal of Economics*, 42(3): 527–554.

Azoulay, P.; Jones, B. F.; Kim, J. D.; and Miranda, J. 2019. Research Efficiency: Turnover, Incentives, and Breakthroughs. *American Economic Review*, 109(7): 2358–2394.

Balija, S. B. 2025. FedMM-X: A Trustworthy Framework for Federated Multi-Modal Learning. *arXiv preprint arXiv:2503.19564*.

Balija, S. B.; et al. 2025a. Decoding Federated Learning: The FedNAM+ Conformal Revolution. *arXiv preprint arXiv:2506.17872*.

Balija, S. B.; et al. 2025b. Fortifying the Agentic Web: A Unified Zero-Trust Architecture Against Logic-layer Threats. *arXiv preprint arXiv:2508.12259*.

Balija, S. B.; et al. 2025c. The Trust Fabric: Decentralized Interoperability for the Agentic Web. *arXiv preprint arXiv:2507.07901*.

Beltagy. 2019. SciBERT: A Pretrained Language Model for Scientific Text. *EMNLP*, 3615–3620.

Bornmann, L.; and Wouters, P. 2021. Artificial Intelligence in Research Evaluation. *Research Evaluation*, 30(3): 1–10.

Butler, K. T.; Davies, D. W.; Cartwright, H.; Isayev, O.; and Walsh, A. 2018. Machine Learning for Molecular and Materials Science. *Nature*, 559: 547–555.

Chi, M.; and VanLehn, K. 2022. Augmented Learning Dynamics in Human-AI Collaboration. *Artificial Intelligence in Education*.

Cortez, P. . 2014. Student Performance [Dataset]. UCI Machine Learning Repository. [Data set].

Dongruo Zhou, Q. G., Lihong Li. 2020. Neural Contextual Bandits with UCB-based Exploration. *arXiv preprint arXiv:1911.04462*.

Erfan. 2025. Elastic MIG Reconfiguration with PCIe-Aware Placement for Multi-Tenant GPUs. *Proceedings of the 11th International Workshop on Serverless Computing*.

Erfan, D.; et al. 2025. Predictable LLM Serving on GPU Clusters. *arXiv preprint arXiv:2508.20274*.

Fortunato, S.; Bergstrom, C. T.; et al. 2018. Science of Science. *Science*, 359(6379).

Gottweis, H.; and Marcus, G. 2025. Towards AI Co-Scientists. *Nature*.

Lattimore. 2020. *Bandit Algorithms*. Cambridge University Press.

Lu, K.; and Zhang, W. 2024. AI-Driven Autonomous Scientific Discovery. *Nature Machine Intelligence*.

Nanda. 2025. FedNAMs: Interpretability Analysis in Federated Learning. *arXiv preprint arXiv:2506.17466*.

Newman. 2004. Coauthorship Networks and Patterns of Scientific Collaboration. *Proceedings of the National Academy of Sciences*, 101(suppl_1): 5200–5205.

Pardos, Z. 2013. Title of the Dataset. [Data set].

Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement Learning: An Introduction*. MIT Press, 2 edition.

Thompson, W. R. 1933. On the Likelihood That One Unknown Probability Exceeds Another. *Biometrika*, 25(3–4): 285–294.

Wang. 2023. Adaptive Experimentation for Science. *Proceedings of the National Academy of Sciences*, 120(12): e2215679120.

Wang, D.; and Barabási, A.-L. 2022. Scientific Discovery in the Age of Artificial Intelligence. *Nature Reviews Physics*, 4(9): 1–14.

Wang, Y.; Liu, X.; and Li, J. 2023. Knowledge Graphs and Graph Neural Networks for Scientific Discovery. *Nature Machine Intelligence*, 5: 1–12.