Integrating Skill Attainment and Enterprise Modeling into Optimal Training Event Scheduling

Eric Haney Ph.D., Rob Harrill, Ryne Spears
Lone Star Analysis
Addison, Texas
ehaney@lone-star.com, rharrill@lone-star.com, rspears@lone-star.com

ABSTRACT

The scheduling of pilot training events is a highly complex and data-intensive process that relies heavily on manual heuristics, often resulting in suboptimal outcomes. Recent advancements in predictive algorithms have demonstrated the ability to forecast student proficiency based on training exposure. Combined with data on current student performance and projected resource availability, the foundational elements exist to enable a dynamic scheduling approach tailored to both individual and enterprise-wide needs. In a representative experiment, an optimal syllabus allocated several hundred decision variables to maximize cohort proficiency in under five minutes – this is referenced to the current paradigm of multi-month processes to reallocate syllabi in most training enterprises.

At the core of this approach is a Skill Attainment Model (SAM) that accounts for the complexity of skills, prior exposure, and the quality and quantity of proposed training events. This model provides a directional estimate of how proficiency will improve with training or degrade due to inactivity. Applications of a SAM in Air Force training command and Naval strike contexts will be discussed.

Complementing this is the Training Enterprise Model (TEM), which represents the capacity of the broader system to support training schedules—incorporating hard assets such as classrooms, aircraft, and runways, alongside human resources like instructors and maintainers. Examples of the TEM from both Air Force and Navy pilot training schools are used for calibration and verification.

Together, these models offer a holistic representation of the training environment. However, the decision space is massively multi-dimensional, and the cadence of real-world training demands automation to avoid overwhelming human schedulers. A non-convex optimization framework is applied to generate automated schedule suggestions that maximize student cohort proficiency without exceeding enterprise capacity or causing unintended ripple effects across the system.

This integrated approach has been applied across both active-duty training pipelines and introductory student syllabi, yielding near-term optimal schedules that reduce instructor workload and improve force readiness. By combining predictive modeling with automated optimization, this method represents a transformative shift in how military training enterprises can balance individual skill development with operational constraints.

ABOUT THE AUTHORS

Eric Haney, Ph.D., is the Chief Technology Officer at Lone Star Analysis. He is responsible for the development, deployment, and support of multiple analytics platforms, including TruNavigator™ and TruPredict™. He also leads the research and development division at Lone Star, Cipher Alchemy. As part of his work, he has been awarded two patents in edge analytics and digital twins. He holds a Ph.D. in Aerospace Engineering (University of Texas at Arlington) and a B.S. in Aerospace Engineering (Texas A&M University).

Rob Harrill is Vice President of Operational Optimization Solutions for Lone Star Analysis and has served in this role since 2017. He is responsible for development and formulation of all optimization solutions for projects and leads interactions with all clients. Mr. Harrill has over 33 years in the defense industry with 26 of that serving on active duty in the US Navy as an Aircraft Maintenance Officer followed by roles in NAVAIR in acquisitions. He was raised in

League City, Texas and graduated with a Bachelor of Science in Aeronautics degree from Embry-Riddle Aeronautical University.

Ryne Spears is a Programs Manager at Lone Star Analysis. He is responsible for production optimization analytics for the U.S. Air Force. Products under management include training pipeline optimization, skill proficiency, and asset lifecycle modeling for various training commands. He holds a B.S. in mechanical engineering (University of North Texas).

Integrating Skill Attainment and Enterprise Modeling into Optimal Training Event Scheduling

Eric Haney, Ph.D., Rob Harrill, Ryne Spears
Lone Star Analysis
Addison, Texas
ehaney@lone-star.com, rharrill@lone-star.com, rspears@lone-star.com

BACKGROUND & MOTIVATION

Schoolhouses responsible for producing qualified pilots operate as highly complex, interdependent systems of systems. The ultimate output of these training enterprises is a pilot who is both qualified—having met the minimum standards—and proficient—capable of executing required tasks effectively and safely. However, determining a "qualified" and "proficient" pilot is itself nontrivial, as different missions, aircraft types, and operational demands set varying thresholds for skill mastery. Moreover, the level of proficiency required is not absolute; it is shaped by the strategic needs of the enterprise and by operational realities—how much proficiency is necessary, and at what cost?

The process of training pilots is further complicated by significant constraints that extend beyond the development of a single student. Training pipelines must simultaneously accommodate many student cohorts, each with overlapping or tangential learning tracks that can interfere with each other's progression. The reliance on expensive, maintenance-intense equipment—such as aircraft, simulators, and classrooms—adds additional resource strain, as does the limited pool of experienced instructors needed to deliver high-quality training. These constraints underscore the need for an approach that can systematically optimize the collective throughput of qualified and proficient pilots while accounting for real-world limitations. Addressing these challenges requires a holistic perspective that considers the dynamic interactions between student skill development, resource availability, and enterprise capacity.

Previous efforts in this field of study include day-to-day flight scheduling which may have even more practical constraints (Jacobs, 2014; Verhoeff, 2015; Foraker, 2021; and Sun, 2025) but does not include the temporal stochastic impacts of present decisions. Efforts to quantify pilot skill attainment are numerous (Nekomoto, 2013; Wojciechowski, 2023; and Koritarov, 2024), however they are very often disconnected from the practical capacity of the enterprise to satisfy proposed interventions at scale.

MODELING ELEMENTS

Complex training enterprises, such as those found in pilot schoolhouses, are characterized by their high degree of interrelatedness and the potential for unintended consequences whenever a change is introduced. A seemingly simple policy adjustment—like altering sortie allocation or shifting an instructor's workload—can ripple through the system in unexpected ways, impacting not only student outcomes but also resource availability and maintenance schedules. To effectively model such enterprises, it is essential to adopt a multi-dimensional, multi-domain approach that captures key performance indicators across the entire system. This includes not only student proficiency metrics but also asset utilization rates, instructor workload, equipment maintenance cycles, and administrative throughput, among others.

Furthermore, policy changes within these enterprises rarely, if ever, satisfy all stakeholders simultaneously. An effective model should not only highlight the aggregate impact of proposed changes but also quantify how different communities—students, instructors, maintainers, leadership—are affected. This quantitative insight is critical in offering decision-makers an opportunity to weigh trade-offs transparently and seek a balance that optimizes overall system performance while acknowledging the constraints and needs of individual stakeholders.

Forecasting future events is inherently probabilistic. There is no way to estimate training capacity multiple years in advance without quantifying the likelihood of variable influences – e.g., weather, compounding coursework, student-to-student variation. These require stochastic mathematic representations of multi-variate relationships. Monte Carlo analysis is one approach for this type of problem and has been detailed in previous publications (Roemerman, 2019; Engel, 2023).

Training Enterprise Modeling

The first element of the integrated approach is the **Training Enterprise Modeling (TEM)** component, which has been successfully applied to multiple pilot training schoolhouses. At its core, the TEM conceptualizes a schoolhouse as the product of **capacity** versus **requirement**, where capacity encompasses the resources available to generate trained pilots, and requirement reflects the operational demand for new pilots. In the context of a pilot training enterprise, capacity is driven by a variety of critical resources. These include the **"plant" or "factory"** components—such as airspace availability, runway access, classrooms, aircraft, and simulators—which are necessary for conducting effective training. Additionally, the instructor workforce—comprising full-time staff, reservists, command staff, and contractors—plays a key role in determining throughput.

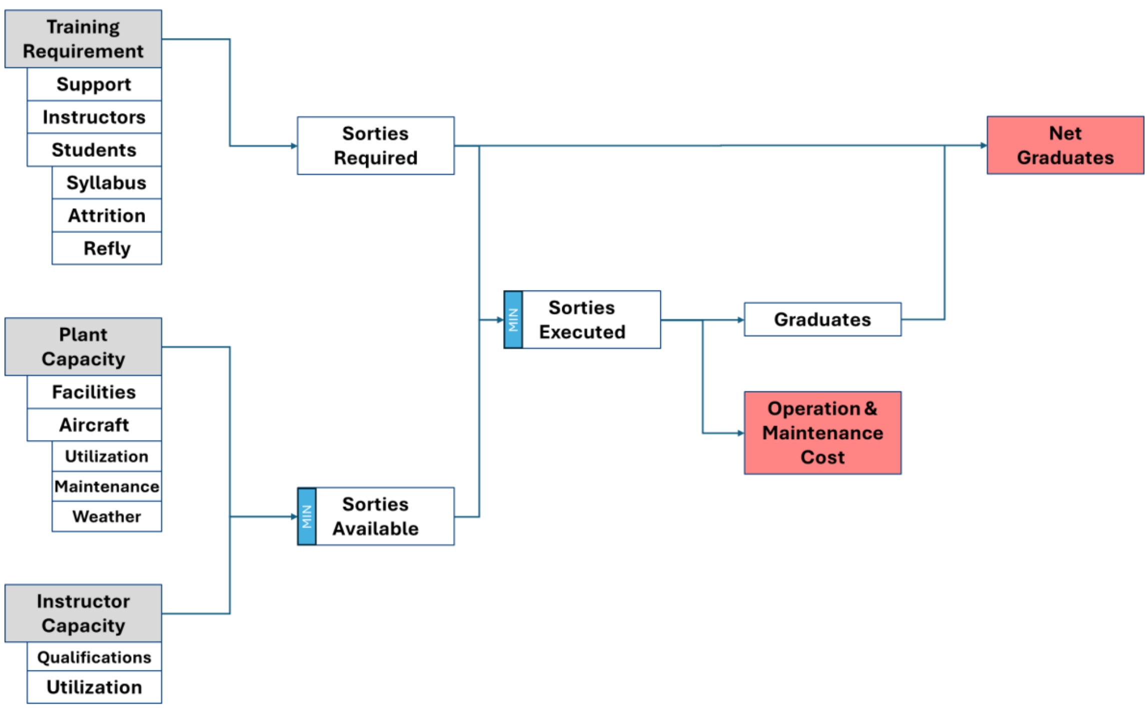

Figure 1. High Level Architecture of Training Enterprise

Conversely, the requirement side of the equation is shaped by operational squadron demands, which dictate how many pilots must be trained to meet force readiness objectives. This is further influenced by factors such as syllabi length, which defines the timeline from initial training to graduation, and the attrition and selection processes, which affect how many students progress through the pipeline. Examples of this modeling approach have been successfully implemented in both the **U.S. Air Force Air Education and Training Command (AETC)** and the **Chief of Naval Air Training (CNATRA)** communities, where it provides a structured way to align training output with dynamic operational requirements.

Two specific outputs have been called out in red within Figure 1 for enterprise optimization. The first is a relative figure of merit called "Net Graduates". This is the difference between the number of graduates expected to be released from the training enterprise and the theoretical maximum number that could be achieved given the same plant and instructor capacity. A Net Graduate value of zero means all resources are being used to the maximum extent possible (although a value of exactly zero leaves little room for process variability or resiliency to change). Values greater than zero represent surplus capacity. Values less than zero mean a deficit is incurred, and therefore less students are graduates than desired. The second figure of merit that was found useful for optimization is the Operations &

Maintenance (O&M) cost – its inclusion allows a degree of fiscal realism to be overlaid on potential scenarios. Both these figures of merit are used in their time-accumulated versions to remove temporal variability and align with standard yearly reporting requirements.

The TEM instance used for further experiments apply this core representation across three training Tracks that use the same pool of devices. Each Track has its own syllabus, student loading, and graduation criteria. This represents a training enterprise that has multiple competing courses which highlights the push and pull of optimal policies. Each Track is then further subdivided into Phases of training that can allow the TEM to be more prescriptive about when certain devices are exposed to the students – i.e., weighted towards classroom time in the first phase, followed by introduction of simulators, transitioning to mostly aircraft. This is a standard representation of typical schoolhouses but could be made more complex with dependent course tracks and multi-site training enterprises. The TEM has been setup with three media types -- Aircraft, Simulator, and Virtual Reality (VR) Headset – to illustrate this tradeoff between availability, cost, and training effectiveness.

Skill Attainment Modeling

The second critical component of the integrated framework is **Skill Attainment Modeling (SAM)**, which forecasts the proficiency of student cohorts as they are exposed to various training events across a spectrum of fidelity and media. Human capacity to learn, absorb, and retain skills is inherently influenced by factors such as time, repetition, and the quality of exposure—elements that any training program must carefully consider. When implementing technical and policy changes, it is essential that the system preserves the output of pilots at the required proficiency levels, despite adjustments in training resources, modalities, or delivery mechanisms.

To represent the progression of student skill development in a systematic way, the model uses a **logistic learning algorithm** that accounts for the complexity of the skill being acquired and the time required to reach an initial novice level. A logistic mathematical representation of proficiency follows the non-linear growth pattern expected in learning complex skills. This structure also allows for flexible calibration, making it adaptable to both well-established skills with abundant training data and new or emerging skills or devices where subject matter expertise may be the only source of insight. Notably, this approach has been successfully applied in **Naval Strike** training contexts as well as in **AETC**'s **Introduction to Fighter Fundamentals**, demonstrating its utility across a range of pilot training pipelines.

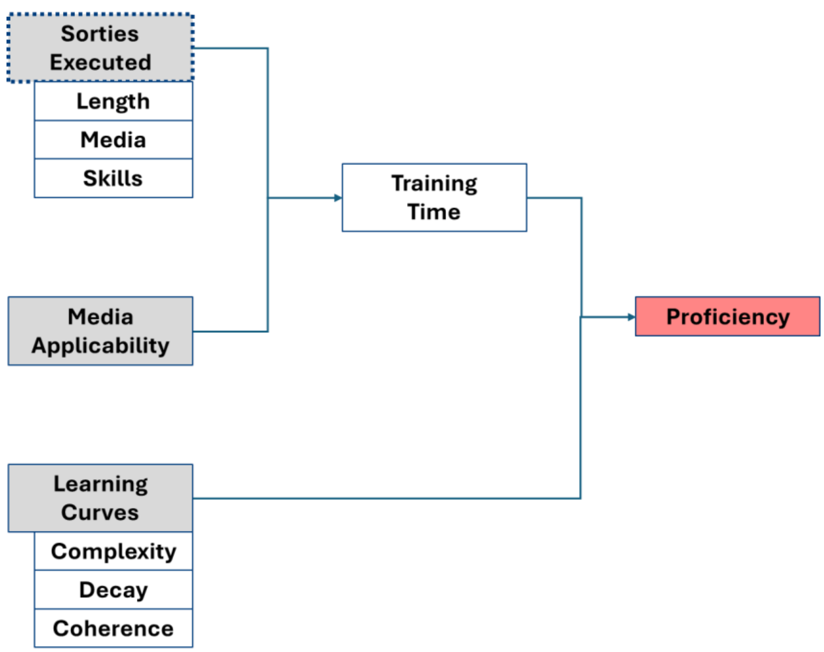

Figure 2. High Level Architecture of Skill Attainment

For the integrated enterprise representation, a SAM reference architecture (shown in Figure 2) was used across 8 skills. Specific skills were not identified in order to keep the model generic of individual schoolhouses – however, the relative relationships between skills and their learning factors mirrored data seen in Air Force Introductory to Fighter Fundamentals (IFF) SAM modeling. Each skill was uniquely perturbed by training exposure across different portions of the training syllabus and by events in different media. This reflects experience with multiple training environments where some skills are better suited to alternative training media and other skills are not addressed until later phases of training. In an integrated enterprise analysis, sorties executed is a translated input from a TEM – a student is only exposed to those skills the enterprise has capacity to execute. There are several extensions to this algorithm that could include: instructional quality to represent variability in feedback, skill decay to represent atrophy of proficiency if not trained consistently, and strict hierarchies of foundational, composite, or dependent skills. These add to the realism of a SAM for some communities but only add needless elaboration for others.

The SAM approach is sensitive to several inputs that may come from a mix of collected field data and elicited subject matter expertise:

- the complexity of learning a specific skill,
- the hierarchical nature of learning advanced skills from the foundation of prerequisite skills,
- the time stability of a skill to inactivity, and
- the effectiveness of different instructional media on learning on a skill.

As the data to each of these inputs is modulated, the skill acquisition curves can vary significantly. Results presented in further discussion within this paper that build off proficiency as a figure of merit, then, could also be modulated by different assumptions or changes in the underlying data to represent actual student skills and training media. The goal is to establish a repeatable mathematical framework that can be replicated in different training communities with their own tailored datasets.

Stochastic Non-convex Optimization

The final component of the integrated approach is a **Stochastic Non-convex Optimization** capability designed to tackle the mathematical challenges of scheduling in complex and uncertain training environments. Unlike simpler problems that can be addressed using deterministic or convex optimization techniques, the pilot training scheduling problem is characterized by discontinuities, objective functions with a mix of narrow local and global minima, and non-linear intermediate relationships that arise from the interaction of multiple resource constraints and training requirements. Such complexity precludes reliance on explicit or well-behaved derivatives, which many classical optimization approaches depend on to guarantee convergence.

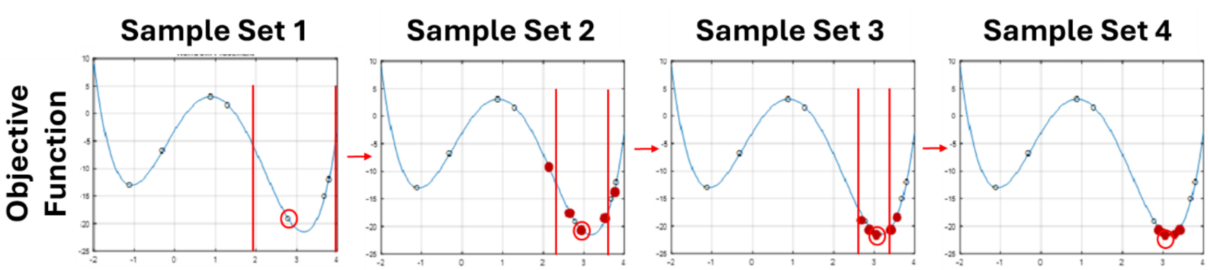

Figure 3. Representative Behavior of Nonconvex Stochastic Optimization

Instead, a stochastic, numerical optimization approach is employed that systematically explores the entire solution space and self-adjusts to identify the most optimal outcomes – illustration in Figure 3 shows the iterative convergence of optimal minima (red points) after randomly sampling the search space (blue points). This approach demonstrates exceptional resilience in converging to effective solutions, even when confronted with large-scale, uncertain, and highly interdependent problems. Notably, this methodology has shown strong results in related defense applications, including **Maintenance Policy Optimization** (Allen, 2021) and **Weapon System Generation Planning** (Allen, 2023). These precedents underscore the robustness of the approach and its suitability for complex training enterprise problems.

HOLISTIC SOLUTION ARCHITECTURE

Representing a holistic training environment requires the integration of both training requirements detailed in a TEM and the student proficiency represented in a SAM. In practice, this could be further extended to maintenance, supply chain, acquisition, or contract management functions as well to increase the breadth of decision-making. As these domains are extended, there is more opportunity to include realistic constraints on the enterprise's decision-making space and in return to limit unintended consequences with a proposed decision – i.e., propose a training regimen that would violate a contract instructors' contract or adoption of a new training device that will not have the production capacity.

Integrating training production with skill proficiency requires linking the Sorties Executed from the TEM and using that interim result to generate training exposure which ultimately influences Proficiency. The mapping of skills that a student is exposed to during a training event within a Phase of a syllabus requires a codified understanding of the specific training command. Both Navy and Air Force pilot training explicitly state the skills that are to be trained by syllabus and specific event, but this mapping can require additional data processing or even elicitation in the case of more freeform training environments. It should be noted that even within rigid environments such as introductory training, exposure to skills in both quantity and quality will always be probabilistic – there are a multitude of variables that will influence exactly how long a student will exercise a targeted skills and similarly, the effectiveness of that exposure to improving proficiency has variable absorption due to latent student variability.

When considering the implementation of mathematical optimization against a stochastic problem with the size and scope of military flight training, caution is required. If given too expansive decision-space, there will exist an infinite number of options and therefore drive no action. If too few constraints are placed upon the optimization, then proposed solutions will suggest invalid, impractical, or detrimental courses of action. The foundational TEM implicitly covers many of this second class of issues by embedding realism boundaries and feasibility checks throughout the calculations. One such representation of available aircraft sorties is shown in the below Figure 4 where each line represents an independent method of calculating daily capacity – blue from aircraft available limited by maintenance and weather, red from flight line operations, and yellow from aggregate yearly fleet utilization. The intensity of the shading represents the degree of uncertainty (commonly P10-P90 and P30-P70 confidence bounds) surrounding the thickest line which is the P50 forecast. The capacity used for downstream calculations is the minimum of these methods. This acts as an implicit constraint that is baked into the foundational model and would therefore be redundant for the sake of optimization.

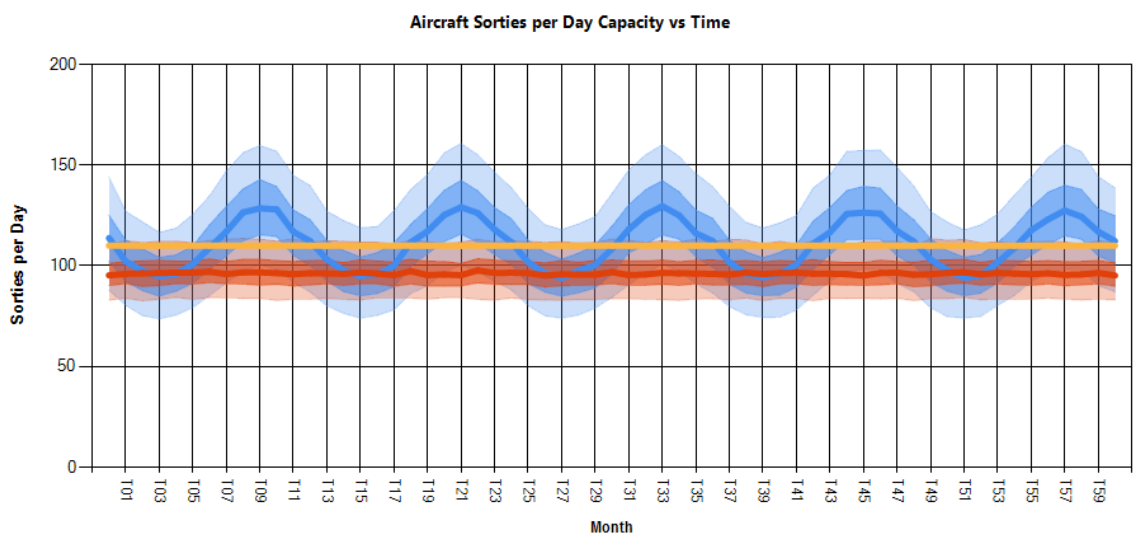

Figure 4. Visual Depiction of Implicit Constraints on Capacity to Train

In contrast, there are several model results that are allowed to fluctuate implicitly and require explicit constraints to produce feasible outcomes. Some examples are shown in Figure 5 below. Cumulative Aircraft Net Graduates is a key constraint that should be kept above 0 – negative values denote there is not enough Aircraft capacity to fully satisfy the production needs. Students are at risk of not graduating within the expected timeframe or will require waivers to progress. Cumulative Operations & Maintenance (O&M) Cost is another potential option for use as an explicit constraint. This could be used to inform budget decisions – i.e., a course of action must be budget neutral. Within the SAM, Overall Enterprise Proficiency is likely to be used as an Objective Function for optimization but could also be used in one of its disaggregated interim formats as an inequality constraint as well. For instance, a scenario could be envisioned that maximizes the proficiency of one cohort by increasing their training regimen while at the same time requiring other cohorts to meet or exceed a required minimum proficiency. This ensures that too many resources are not directed at the target cohort at the detriment of other students. Proficiency at the individual skill level could also be used as a constraint to avoid an overreliance placed on training media or skills that are acute and do not produce a well-rounded graduate.

A helpful exercise before beginning mathematical optimization is to parametrically perturb the system to better understand its behavior. Figure 5 below shows the behavior of potential Constraints (top) and Objective Function (bottom) as required syllabus events are linearly increased from their baseline values on the left to baseline + 10 events for each media type on the right. The top line shows how increased event requirements impacts in both Aircraft Net Graduates and Simulator Net Graduates – it should be noted that the intercept of where these bands cross 0 is where the policy starts becoming infeasible. Cumulative O&M Costs is also shown on the top row and depicts the blended linear relationship between increased events in different training devices on enterprise cost. Overall Enterprise Proficiency (on the bottom) row shows the relative linear relationships between events and blended proficiency across training tracks and phases – it should be noted with different skill input data this relationship adjusts dramatically with the parametric logistic definition.

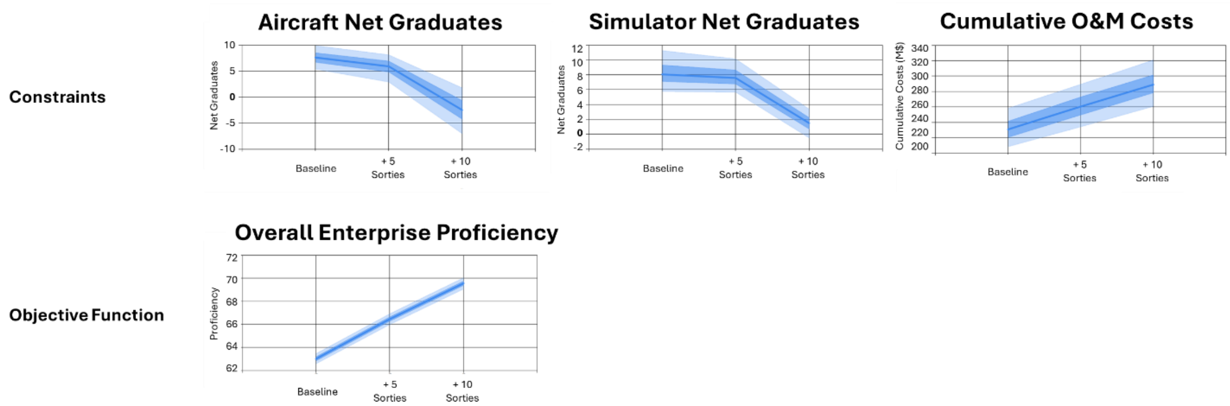

Figure 5. Behavior of Select Enterprise Outputs to Monotonic Increase in Syllabus Events

EXPERIMENTS

The practical illustration of an integrated enterprise optimization will take the same base model architecture but apply different decision scenarios. This may include adjusting the baseline data to represent unique starting points for the decision, as well as modular combinations of constraints and objective functions to represent alternative goals of the proposed policy change.

Fixed-Resource Optimization of a Single Training Track

The first optimal policy experiment is to: *create the most highly proficient pilots graduating from Track A with the given workforce and material resources*. This exercise is aligned with the practical aim of prioritizing a graduating mission set that is likely to be in high demand and budget constraints are removed. For this policy outcome, "Overall Enterprise Proficiency" is used as the Objective Function (although an upstream output for only Track A cohort proficiency may be more semantically consistent). Inequality constraints were set as requiring positive Cumulative Net Graduate at the end of the time-period of interest – for illustration purposes, 12 months. The perturbing variables to optimize are the number of required Syllabus Events for Track A, with allowable variation for each of the 12-month segments.

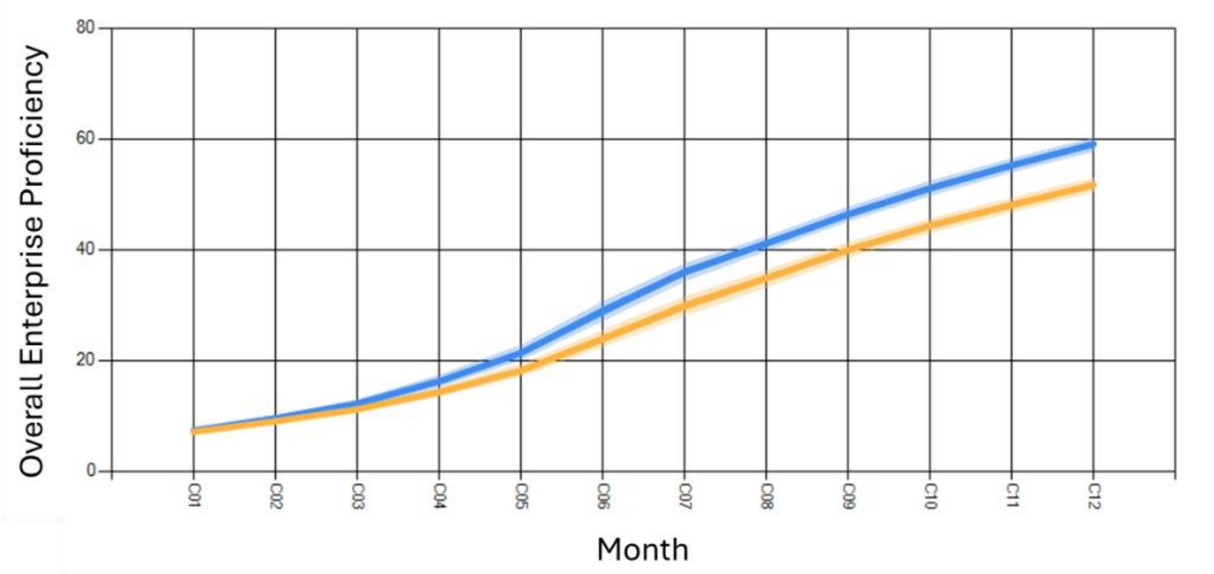

Figure 6. Proficiency over a 12-Month Period in the Baseline (yellow / bottom) and Track A Optimized Policy (blue / top)

Figure 6 above shows the impact to Overall Enterprise Proficiency when the optimal policy is calculated. It shows a noticeable increase in graduate proficiency, accumulating an approximate 10% improvement. It should be noted that the maximum upside of proficiency gain is implicitly restricted in two manners: 1) the SAM limits proficiency values between 1-100 for standardized scaling, and 2) optimized policies are only allowed within prescribed bounds of the stochastically optimized variables. For an example on point two, the baseline number of Syllabus Events required for Track A Phase 1a was 40 Events – the optimal policy was provided a search space of [40,50] to require at least as many events but allow for meaningful increase in events. Increasing this search space for optimization exploration may result in higher proficiency values (i.e., a more optimal outcome) but comes at the cost of higher computational attempts for solution points that are not likely to satisfy the constraints – i.e., Syllabus Events in the aircraft over 50 will often violate explicit constraints for Net Aircraft Graduates and are therefore not realistic points to sample. The manual parametric exploration in Figure 5 illustrated that more than 50 events were very likely to push Aircraft Net Graduates less than 0 and is therefore unlikely to yield feasible results.

Figure 7 shows the proposed policy changes to the Track A syllabus. Each phase of the training track has increased events; however, the magnitude of increase is a function of training device media fidelity and their corresponding events, the portion of the proficiency-training exposure curve, and the excess capacity of each training device. Simulator events do not show the same level of increase (from 13 to 17 nominally) as Aircraft (40 → 47) and VR (30 → 36) training – this highlights the data differences between media types and highlighted skills.

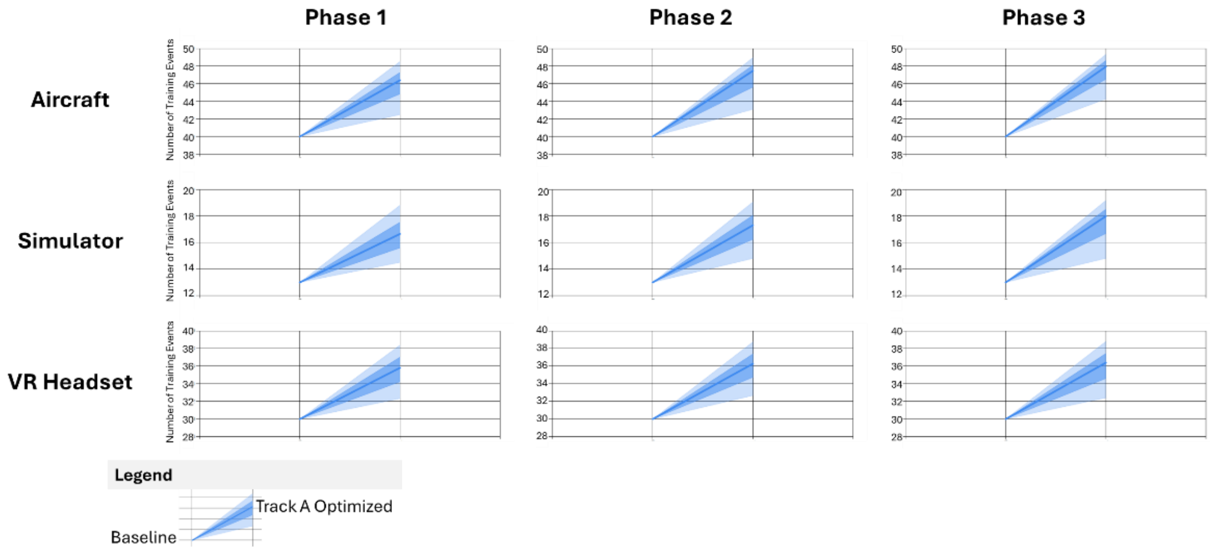

Figure 7. Suggested Syllabus Policy Impact for Optimizing Track A Resource Utilization

Budget-Neutral Syllabus Reallocation

The second optimal policy both widens the search space and applies additional constraints: *given a fixed budget and wide latitude to change syllabus mix by media, produce the most proficient student cohorts.* In this scenario, all three training Tracks and each of their three corresponding Phases is subject to change. Whereas the first experiment required Syllabus Events to be at the previous baseline or higher, the second policy experiment allows events to be removed from media as well. The objective function is still to produce the highest level of Overall Enterprise Proficiency. In contrast, the budget is required to be at or below the current spending levels.

Table 1 shows the suggested optimal policy for each Track / Phase / Media combination. In general, the optimal solution proposes to increase the use of VR Headsets, especially in early phases of training, keep Aircraft usage steady, and balance out cost by removing traditional simulator events. As with the first experiment, Simulator events are less emphasized due to their less significant connection to relevant skills. They are even less emphasized in this second experiment to their relatively higher cost than the VR Headset training devices.

One oddity of stochastic optimization is highlighted in Table 1 – most events were suggested to decrease to keep the overall yearly budget flat. However, this is counterintuitive. It would be expected that as many Track / Phase / Media combinations are decreased, a rough corresponding number of combinations should increase (with some variance due to differences in operating costs between media). The phenomenon highlights that the stochastic optimization algorithm is rejecting policies that could violate either the Net Graduates resource or O&M cost constraints, even if that outcome is very remote. In other words, the prescribed policy is overly conservative relative to constraints and could be relaxed to a more permissive risk tolerance. More Syllabus Events could therefore be allowed with a greater than zero (but small) chance of violating the constraints. These unlikely scenarios could be overcome with active management of the process.

Table 1. Suggested Syllabus Policy Impact of Broad Reallocation within a Fixed Budget

		Aircraft	Simulator	VR Headset
Track A	Phase 1	-1%	-27%	4%
	Phase 2	-10%	-38%	-4%
	Phase 3	-14%	-50%	-9%
Track B	Phase 1	0%	-36%	3%
	Phase 2	-7%	-35%	-1%
	Phase 3	-6%	-59%	-5%
Track C	Phase 1	-3%	-41%	-5%
	Phase 2	-9%	-37%	-5%
	Phase 3	-7%	-60%	-12%

CONCLUSIONS

This work demonstrates that mathematical optimization of large, complex, and uncertain enterprises is not only technically feasible but also holds significant promise for streamlining decision-making in domains with high operational tempo and resource constraints. Military flight training presents a compelling use case for such optimization due to its dynamic environment, complex scheduling requirements, and the high-stakes nature of its outcomes.

Optimal syllabus design was illustrated for a complex training enterprise within 5 minutes on a commercial laptop – this is in comparison to a typical multi-month manual process for most schoolhouses. There is obviously nuance into the finer points of training event details but rapidly producing an optimal starting point for detailed discussions or experimentation (such as the U.S. Air Force's use of small test groups of student pilots) provides massive return on analytics. These improvements unlock enterprise training managers to conceive new paradigms for course design – e.g., more responsive incremental changes to reflect student response, proficiency emphasis, or media availability. In a defense environment with rapidly changing technology and geopolitical concerns, being able to pivot large institutions will be a key enabler.

The same representation of a training enterprise can be re-purposed for additional optimal policy generation by selecting different objective functions, constraints, and decision boundaries. Ongoing development should focus on identifying high-value, repeatable optimization use cases that align with stakeholder needs. These use cases can be transformed into parameterized workflows that allow for the rapid generation—and re-generation—of tailored training approaches. Each of these provide step functions in saved time, manpower, and resources to adjust training policies. Some examples of optimal decision policy that would greatly improve enterprise schoolhouse decision outcomes include:

- Optimal instructor allocation across a geographically-dispersed training enterprise
- Optimal assignment of scarce media when acquiring new training assets
- Optimal student inductions to balance resource utilization, student inactive time, and graduated pilots

Additionally, this modeling framework can be scaled both upward to address enterprise-wide decisions and downward to focus on individual student optimization. This multi-scale application will allow a diverse range of stakeholders—strategic planners, schedulers, and instructors—to engage with the model at a level appropriate to their decision-making horizon.

One key insight from this effort is that transitioning from an analyst-in-the-loop approach to a mathematically optimized model requires careful and deliberate definition of decision variables, constraints, and objective functions. In many cases, the foundational training model itself may need to be revised to integrate effectively with an optimization framework. Incorporating constraints from multiple domains—such as resource availability, instructor limits, and student progression—adds realism and increases the operational value of the resulting course of action.

However, excessive rejection of sample solutions due to constraint violations can inadvertently reduce the search space and limit feasible outcomes, highlighting the need for a balanced and flexible optimization structure.

Both military and commercial training enterprises face growing pressure to produce qualified, proficient pilots in greater numbers and shorter timeframes—all while operating within constrained budgets and limited instructor availability. The global shortage of skilled aviators has driven an urgent search for technology-enabled solutions that can improve the efficiency and effectiveness of pilot production. The integrated training enterprise modeling and optimization framework presented here offers a promising avenue for meeting this demand. By aligning resource usage with student performance and operational requirements, it has the potential to significantly enhance the capability of flight training schoolhouses across both sectors.

ACKNOWLEDGEMENTS

A special thanks to Dr. Randy Allen for his contributions to nonconvex stochastic optimization that were foundational to this work – congratulations on the patent award!

REFERENCES

Allen, R. Evolved AI® for First-Order Conceptual Missile Design Optimization and Threat Assessment. (2023). 2023 AIAA Defense Forum.

Allen, R. (2023). Hypersonic Missile Thread Modeling, Simulation, and Assessment. (2023). 2023 Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC paper 23146).

Allen, R., Engel Z., and Haney E. Evolved AI® for Model-Based Reinforcement Learning. 2021 Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC paper 21199).

Engel, Z., Harrill, N., Velez, N., Ediger, J., & Allen, R. (2023). Continuous Asymmetric Risk Analysis (CARA). 2023 Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC paper 23270).

Foraker, J., Lazzaro, G., Nelson, P., & Barkley, P. (2021). Scheduling of Daily Flight Training for a United States Navy Strike Fighter Squadron Detachment. *Military Operations Research*, *26*(2), 5-24.

Jacobs, R. S. (2014). *Optimization of daily flight training schedules* (Doctoral dissertation, Monterey, California: Naval Postgraduate School).

Koritarov, T. (2024). ADVANCING UNMANNED AERIAL VEHICLES PILOTS TRAINING: INTERACTIVE TECHNOLOGIES AND GAMIFIED LEARNING FOR ENHANCED SKILLS ACQUISITION. *Norwegian Journal of development of the International Science No*, *144*, 144.

Nekomoto, T. Y. (2013). A systems analysis of strike Naval Aviation training.

Sun, P., Liu, J., & Nian, H. (2025). An algorithm for automatically arranging flight training plans. *Bulletin of the Polish Academy of Sciences Technical Sciences*, e151672-e151672.

Roemerman, S., Volpi, J., & Allen, R. (2019). Executive Risk Assessments for the Age of Algorithms. 2019 Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC paper 19110).

Verhoeff, M., Verhagen, W. J. C., & Curran, R. (2015). Maximizing operational readiness in military aviation by optimizing flight and maintenance planning. *Transportation Research Procedia*, *10*, 941-950.

Wojciechowski, P., Wojtowicz, K., & Błaszczyk, J. (2023, November). AI-driven method for UAV Pilot Training Process Optimization in a Virtual Environment. In *2023 IEEE International Workshop on Technologies for Defense and Security (TechDefense)* (pp. 240-244). IEEE.