Proceedings of the Annual General Donald R. Keith Memorial Conference
West Point, New York, USA
April 23, 2026
A Regional Conference of the Society for Industrial and Systems Engineering

Monte Carlo Optimization of Snowmaking Operations at Keystone Resort

Aden Extrand, Raymond Picard, Jared Switzer, Jacob Wielhouwer, Gerry Gonzalez, Taylor Leonard, and Benjamin Kallemyn

Operations Research Program, United States Air Force Academy, Colorado Springs, Colorado 80841

Corresponding author's Email: gerry.gonzalez@afacademy.af.edu

Author Note: Aden Extrand, Raymond Picard, Jared Switzer and Jacob Wielhouwer are first-class cadets at the U.S. Air Force Academy (USAFA) majoring in Operations Research and Gerry Gonzalez, Taylor Leonard, and Benjamin Kallemyn are faculty advisors. The views expressed herein are those of the authors and do not reflect the position of the United States Air Force Academy, the Department of the Air Force, or the Department of War.

Abstract: We develop a physics-based Monte Carlo simulation to optimize snowmaking for progressive terrain openings at Keystone Resort, Colorado. The model predicts opening dates by integrating snow-gun performance curves, elevation-based thermodynamics, capacity-constrained resource allocation, and dynamic snowpack processes. We calibrate the simulation over five seasons, achieving a mean absolute error of 104.8 hours. A 1,000-trial block-bootstrap evaluates configurations and operating policies at identified bottlenecks. Results demonstrate that an Inverted configuration involving the reallocation of fan guns to high-elevation terrain, combined with a relaxed dependency policy advances terrain access by up to 110 hours, while suboptimal equipment choices may delay openings by 249 hours. Statistical tests indicate large effect sizes with high significance. Our framework provides actionable guidance for resort leadership to prioritize infrastructure investments and operating policies that accelerate terrain availability, maximize early-season revenue, and reduce seasonal operational variability.

Keywords: snowmaking modeling, Monte Carlo simulation, ski resort operations, snow physics

1. Introduction

1.1 Background and Motivation

Snowmaking infrastructure enables Keystone Resort to provide dependable early-season skiing and maintain terrain during periods of limited natural snowfall. Earlier terrain openings allow the resort to open restaurants and other on-mountain services sooner, generating additional revenue. As a result, early-season terrain availability creates tangible operational and financial value. To support these operations, Keystone maintains an extensive snowmaking system that includes snow guns, pipelines, pumps, compressors, and weather stations that deliver pressurized water and compressed air to trails during viable weather windows. Resort leadership prioritizes capital investments that improve early-season reliability and operational efficiency.

1.2 Problem Statement

Keystone Resort requires a quantitative framework to evaluate the operational impact of snow-gun upgrades. The complex interaction of water supply, compressed air, and weather creates system bottlenecks that frequently limit production. Current investment decisions rely on operator experience and management judgment to identify priority infrastructure improvements. As a result, it remains difficult to determine whether proposed investments in new snow guns and infrastructure meaningfully improve early-season terrain availability given existing system realities. The resulting model quantifies how snow-gun upgrades translate into operational gains under realistic capacity limits and weather conditions.

1.3 Related Work

Our study integrates snow physics, operations research, and uncertainty analysis to develop a robust decision-support tool. The snowpack model incorporates a multi-stage mass balance of fresh-snow density following Hedstrom and Pomeroy (1998), while long-term evolution uses Anderson's (1976) exponential densification. We model snow loss using temperature-index melt processes from Lehning et al. (2002). An iterative solver calculates wet-bulb temperature, the

thermodynamic limit of evaporative cooling, using Ferrel's (1886) psychrometric constant and Murray's (1967) saturation vapor pressure refinement. We adjust atmospheric pressure for elevation per the U.S. Standard Atmosphere (1976). The framework applies capacity-constrained proportional logic to staged terrain openings (Hanzer et al., 2020; Spandre et al., 2018). We represent meteorological uncertainty with bootstrap sampling (Efron, 1979), employing a block bootstrap approach to preserve autocorrelation inherent in weather patterns. Calibration utilizes differential evolution (Storn & Price, 1997) and local refinement (Byrd et al., 1995) to optimize the non-linear snowmaking parameter space. Finally, we evaluate results using Mann-Whitney (1947) significance and Cohen's (1988) effect size thresholds.

2. Data Description

Our study integrates publicly available datasets with proprietary information provided by Keystone Resort for research purposes. Hourly meteorological data from October 2005 to February 2026 were sourced with the Open-Meteo API, capturing air temperature, dew point, relative humidity, solar radiation, wind speed, and snowfall to model seasonal snowmaking conditions. Historical terrain-opening records for the 2020–21 through 2024–25 seasons, compiled from archived resort reports and press releases, document individual run openings and provide empirical reference points for early-season availability. The available records restrict our training and test data to a limited sample of 5 snowmaking seasons. We digitize spatial trail boundaries in Quantum Geographic Information System (QGIS) using public trail maps and a phase-prioritized polygon file provided by Keystone. Elevation data are derived from 1-meter digital elevation models (DEMs) sourced from United States Geological Survey (USGS) LiDAR products. Additionally, we georeference snow gun locations using resort infrastructure map overlays to assign precise coordinates and elevation attributes to each unit. System-wide water and compressed air capacities are quantified by cross-referencing visible pump house asset identifiers with manufacturer performance specifications, converting aggregate horsepower ratings into operational flow limits with standard industrial coefficients. All datasets are stored in standard CSV and GeoPackage formats to support transparent access and reproducibility.

3. Model and Simulation Methodology

3.1 System Description

Keystone Resort organizes early-season snowmaking around seven prioritized terrain phases (Phases 1–5, 7, and 8). Our model excludes Phase 6 because its back-mountain location remains physically inaccessible to guests until the primary transit corridor of Phases 5 and 8 are operational. This is consistent with current Keystone leadership prioritization. Each phase represents a geographic area with defined trails, elevations, and installed snowmaking infrastructure. Terrain connectivity creates key operational dependencies. As seen in Figure 1, Phase 8 (pink) must open before Phase 5 (dark green) because it provides the skiable route required to reach that portion of the mountain. Phase 5 also enables access to Outpost Plaza (star), an on-mountain restaurant/shop, so delays in Phase 8 directly delay Phase 5 and restrict guest services. Phase 7 (blue) operates independently due to its separate location and terrain connectivity. All points in Figure 1 indicate snow gun infrastructure.

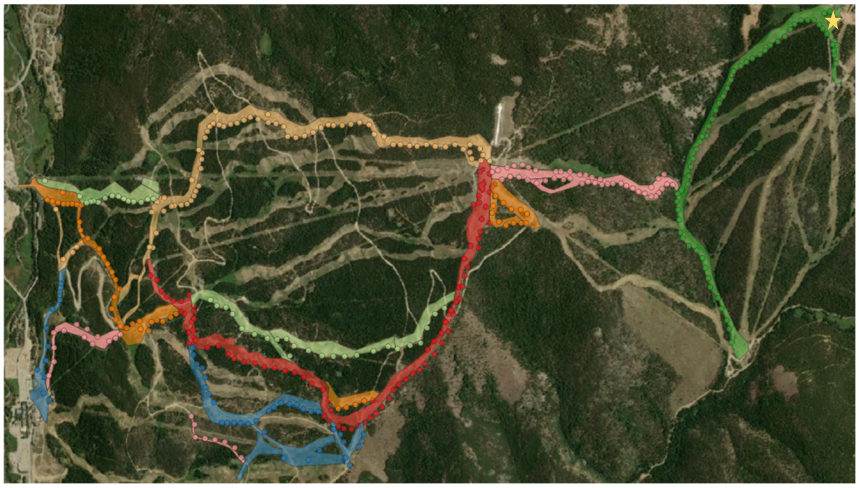

Figure 1. Full Mountain QGIS Mapping

3.2 Model Components and Calibration

3.2.1 Weather and Natural Snow

Our model interpolates weather data from four stations (RR, CB1, Summit, and Bergman) with elevation-based linear weighting. A solar-adjustment factor modifies air temperature using LiDAR-derived slope and aspect to account for radiation on south-facing terrain. We compute wet-bulb temperature using an iterative psychrometric solver adjusted for atmospheric pressure to satisfy ambient vapor pressure with evaporative cooling. Snowmaking operates below a 27 °F wet-bulb threshold with production efficiency increasing significantly as temperatures drop. The algorithm converts natural snowfall to snow-water equivalent (SWE) using a fresh-snow density model for freezing conditions and linear approximations for warmer intervals. Combined natural and machine-made SWE contribute to a dynamic snowpack that evolves through exponential densification and temperature-index melt processes.

3.2.2 Snow Gun Types and Artificial Snow

Our analysis utilizes five snow gun types to operate across the terrain phases based on Keystone Resort's existing infrastructure inventory and operational data provided for this analysis: the Borax 6M Snow Gun, a high-output air-water gun; the K-Gun Snow Gun, a legacy mobile unit used exclusively in Phase 5; the York 20 Snow Gun, a compressed air stick gun; the TechnoAlpin TF10-3M, an electrically powered fan gun that does not require compressed air; and the TechnoAlpin TF10-6M, a tower-mounted fan gun that improves snow making and distribution efficiency. We distribute snow guns across terrain phases with known locations and elevations, which determines the spatial distribution of production capacity. Each gun type includes performance curves that specify gallons per minute (GPM) and cubic feet per minute (CFM) as functions of wet-bulb temperature.

3.2.3 Snowpack Physics

Our model treats the snowpack as a dynamic mass balance where total snow depth is the sum of machine-made and natural SWE components divided by their respective densities for phase i at hourly timestep t in inches:

$$Snow\ Depth_{i,t} = \frac{SWE_{machine_{i,t}}}{\rho_{machine_i}} + \frac{SWE_{natural_{i,t}}}{\rho_{natural_t}} \quad (1)$$

We hold machine-made snow density at a calibrated constant to reflect immediate compaction from grooming. In contrast, natural snow density evolves through exponential densification:

$$\rho_{natural_{i,t}} = \rho_0 + (\rho_{max} - \rho_0) \times (1 - e^{-kt}) \quad (2)$$

where ρ_0 is initial fresh-snow density, ρ_{max} is maximum potential density, and k is a settlement coefficient. When hourly natural SWE is added to the pack, we recalculate the bulk natural density to preserve mass-volume consistency.

3.2.4 Resource Allocation

At each hourly timestep, we calculate total water and compressed-air demand by summing the requirements of all operating snow guns. We enforce feasibility by comparing instantaneous demand to system capacity limits that the river pump station and downstream booster stations define, which maintain pressure rather than adding capacity. To manage these constraints, we utilize a two-stage logic. First, we assign a system reserve share of 0.05 to all active phases to prevent total production stalls. Second, we redistribute the remaining water and air capacity according to a priority scheme. In this stage, we assign Early Season (Phases 1–5), Phase 7, and Phase 8 specific weights that govern the proportional scaling and redistribution of surplus capacity. This framework ensures aggregate water and air use stay within system limits while prioritizing available resources across active units according to resort-defined terrain opening strategies.

3.2.5 Accumulation Simulator

Our accumulation simulator functions as a discrete-time model that builds the snowpack through an iterative, hourly process. For each terrain phase, we determine the total snow depth in inches at the conclusion of the simulation window T by summing hourly machine-made and natural snow contributions while adjusting for cumulative mass loss:

$$Final\ Snow\ Depth_{i,T} = \sum_{t=0}^{T} \left[\frac{\eta_i \times Production_{i,t}}{\rho_{machine_i}} + \frac{\Delta SWE_{natural_{i,t}}}{\rho_{natural_{i,t}}} \right] - Melt_{i,T} \quad (3)$$

In Equation 3, $Production_{i,t}$ represents the hourly water volume delivered to the phase, which is converted to SWE with the water-to-snow efficiency factor η_i. The first term within the summation accounts for the machine-made contribution to total

depth, while the second term captures natural accumulation adjusted for dynamic natural snow density. We correct final snow depth by removing $Melt_{i,T}$, which represents the snow lost to weather conditions over the duration of the simulation.

From October 1 through December 31, the simulator updates weather interpolation across stations at each hourly time step, evaluates individual snow-gun performance curves, executes resource allocation logic, and updates the snowpack mass balance. The simulation for a specific phase terminates once the hourly depth reaches the 18-inch target threshold. Regardless of the specific hour we achieve this threshold, the model logs the effective terrain opening at 0900 local time on the following calendar day to align with ski patrol operating windows. Key recorded metrics include the opening date, total water consumption, and capacity-binding frequency, defined as the percentage of eligible hours where system demand exceeded available capacity. We implement the simulator in Python, utilizing NumPy for vectorized operations and Pandas for time-series management.

3.2.6 Data and Metrics

We use historical terrain opening records from the 2020–21 through 2024–25 seasons for model calibration. We derive observed opening dates for each terrain phase from snowmaking logs, using the earliest and latest recorded snowmaking dates within each phase to define operational start and validation opening times. For each simulated season, our model predicts a terrain opening when simulated snow depth exceeds the 18-inch operational threshold. We evaluate model performance by comparing predicted and observed opening times using mean absolute error (MAE) and bias in hours. We use a combined score as the optimization objective to penalize systematic timing errors while prioritizing overall accuracy:

$$Score = MAE + 0.5 \times |\ Bias\ | \tag{4}$$

3.2.7 Three-Stage Optimization

We calibrate model parameters through a three-stage optimization process using a differential evolution algorithm followed by L-BFGS-B local refinement. Stage 1 estimates global resource allocation parameters, including the system reserve share and priority weights for early-season terrain (Phases 1–5), Phase 7, and Phase 8. Stage 2 calibrates phase-level snowpack and production parameters. We optimize Phases 1 and 2 jointly as they open together as the first available terrain, while optimizing Phases 3, 4, 5, 7, and 8 independently. Stage 3 executes a final joint optimization across all parameters to minimize the combined error between simulated and observed terrain opening times.

3.3 Monte Carlo Simulation

3.3.1 Baseline and Scenarios

The baseline configuration for Phase 5 consists of 58 snowmaking guns: 39 York air-water guns on Paymaster (upper mountain) and 19 TF-10 3m water-only fan guns on Mozart Flats (lower mountain). While gun locations and elevations remain fixed, we evaluate three scenarios by altering Phase 5 gun types. The All Fan configuration uses 58 TF-10 3m units, representing a best-case upgrade that requires 39 additional fan guns and electrical infrastructure on Paymaster. The All York configuration uses 58 York units as an inefficient worst-case for validation and requires the purchase of 19 York guns. The Inverted configuration places 39 TF-10 3m units on Paymaster and 19 York units on Mozart Flats, which requires 20 new fan guns and the installation of power on Paymaster. We apply each scenario by updating Phase 5 CSV snow gun records while preserving geometry to compute opening times relative to the baseline.

3.3.2 Monte Carlo and Dependency Design

We evaluate each configuration using 1,000 trials. Each trial constructs a synthetic season from October 1 to January 15 using block-bootstrap resampling. We draw consecutive 10-day blocks from a multiyear pool to capture meteorological variability and preserve the autocorrelation of warm and cold temperature spans. The model applies these seasons to the baseline and all Phase 5 scenarios. Operating hourly, the system enforces eligibility with a minimum 4-hour block at or below the wet-bulb threshold of 27 °F. A two-stage policy allocates water and air capacity through reserved shares and weight-based redistribution. The model converts water to machine-made SWE using phase-specific efficiency, adds natural SWE through temperature-dependent density, and updates pack density through exponential densification. After applying degree-hour melt, the model declares a phase open when density-aware depth reaches 18 inches. We then map raw opening times to effective openings based on the dependency policy. The Strict policy requires Phase 8 to open before Phase 5, whereas the Relaxed policy allows Phase 5 to open independently. The all-runs open time represents the maximum of effective openings across all phases.

3.3.3 Simulation Statistical Analysis

For each trial and scenario, the model records Phase 5 opening times under both Strict and Relaxed policies along with the all-runs-open time. We compute delta hours relative to the baseline and summarize these distributions using medians, percentiles, means, and standard deviations. To compare Strict versus Relaxed outcomes, we apply Welch's t-tests and Mann-

Whitney U-tests to the delta hour distributions. Our analysis reports Cohen's effect sizes and 95 percent confidence intervals to interpret differences in means. For the three Phase 5 configurations, Cohen's d values of 1.566, -1.481, and 1.631 indicate large effect sizes, which demonstrates that the dependency assumption substantially alters model outcomes. Conversely, Cohen's d values of 0.000 for the All Runs Open simulations indicate no significant effects. We assess statistical significance at the 0.05 and 0.01 levels. Summary statistics for each scenario include the proportion of trials opening at least 24 hours, 7 days, and 14 days earlier than the baseline.

4. Results and Analysis

4.1 Results

Calibration of our model over 2020–21 to 2024–25 yields a weighted MAE of 104.8 hours (4.1 percent of the 107-day snowmaking simulation window) and bias of 5.0 hours, indicating strong model precision. Key parameters include a 0.05 system reserve share and priority weights of 3.21 (early-season), 0.52 (Phase 7), and 0.53 (Phase 8). Phase-specific values include water-to-snow efficiency 0.41–0.59, snow density 471–554 kg/m³, melt rate 0.08–0.18 in/°F-day, area fraction 0.59–0.82.

Table 1. Model Calibration Performance Metrics by Season

Season	Type	MAE (hours)	Bias (hours)	Score
2020–21	Calibration	89.7	39.4	109.4
2021–22	Calibration	95.3	44.1	131.1
2022–23	Calibration	57.7	23.4	69.4
2023–24	Calibration	123.6	-99.3	173.2
2024–25	Test	157.7	17.1	166.3
Overall	**Weighted Average**	**104.8**	**5.0**	**107.3**

Model precision (Table 1) fluctuates with seasonal climate variability. The 2022–23 season presents the highest accuracy during the coldest winter since 2013 and the 11th wettest on record. Conversely, 2023–24 season exhibits the highest error training, driven by the 7th warmest December on record and a 10th percentile snowpack. Our 2024–25 test run is similarly impacted by the 2nd warmest and top 10 dryest October in 130 years, which severely restricts early-season snowmaking windows.

Table 2. Monte Carlo Simulation Results: Median Change in Opening Time Relative to Baseline

Scenario	Phase 5 Opening: Strict (hours)	Phase 5 Opening: Relaxed (hours)	All Runs Open
All York	+1	+249	-6
All Fan	-61	-125	0
Inverted	-15	-110	-2

The negative values in our Monte Carlo results (Table 2) show that fan guns consistently accelerate Phase 5 openings, especially with relaxed dependencies, while York guns delay phase access. Observed results show that snow gun type changes have minimal effect on all-runs-open timing, confirming Phase 8 as the bottleneck. Statistical tests confirm large, significant differences between dependency policies (Cohen's $d > 1.5$, $p < 0.001$).

4.2 Analysis

We assess model robustness through two complementary tests. First, we perturb key calibrated parameters within operationally plausible ranges (efficiency ± 0.10, machine-snow density ± 50 kg/m³, melt rate ± 0.02 in/°F-day) and rerun the model using 100 Monte Carlo trials per perturbation. While a reduction from the primary 1,000-trial experimental design, 100

trials provides a sufficient computational trade-off to confirm trend stability. Across all test variations, configuration performance remained consistent (All Fan earliest, Inverted second, All York latest), and Phase 5 median effective opening dates shift by at most 48–72 hours. Density perturbations produce the widest spread, while efficiency and melt rate produce smaller shifts. Second, we evaluate simulation-level sensitivity by applying the same parameter perturbations to 100-trial subsets under both Strict and Relaxed dependency policies. All-runs-open times remain stable with median changes at or below 4.8 hours across all settings. The same configuration performance pattern held under both policies, with smaller effect magnitudes under Strict. These results confirm that our qualitative conclusions and operational recommendations are robust to reasonable uncertainty in calibration inputs, with Phase 5 uncertainty bands of 48–72 hours on scenario medians.

5. Conclusions and Recommendations

Our study presents a simulation-based framework for optimizing snowmaking operations at Keystone Resort, with a focus on accelerating early-season terrain openings. We show that Phase 5 can open, up to 110 hours earlier, with few added resources. Our scenario analysis demonstrates that the most cost-efficient approach is to reallocate fan guns on Paymaster and York guns on Mozart. Implementing this inverted configuration requires the purchase of 20 additional fan guns and installation of electrical infrastructure on Paymaster, leveraging existing assets while avoiding the expense of a complete fan gun upgrade. This solution directly addresses the system's primary operational constraint, the limited capacity of the compressed air supply. By relocating fan guns to Paymaster, additional air capacity becomes available for the York guns on Mozart, enabling them to sustain higher output and take full advantage of favorable weather windows. We expect the redeployment of 20 York guns to Phase 8, replacing legacy K-guns, to further enhance snowmaking efficiency and resource utilization at a critical access point. Adopting a relaxed dependency policy, which allows Phase 5 to open independently of Phase 8, expedites guest access to additional terrain and the Outpost Plaza, and supports the objective of maximizing early-season operations.

6. References

Anderson, E. A. (1976). *A point energy and mass balance model of a snow cover* (NWS 19). National Weather Service, NOAA.

Byrd, R. H., Lu, P., Nocedal, J., & Zhu, C. (1995). A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing*, *16*(5), 1190–1208. https://doi.org/10.1137

Cohen, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd ed.). Lawrence Erlbaum Associates.

Colorado Climate Center. (2022–2024). *Colorado monthly climate summary [Series of monthly reports]*. Colorado State University. https://climate.colostate.edu/monthly_summary.html

Efron, B. (1979). Bootstrap Methods: Another Look at the Jackknife. *The Annals of Statistics*, *7*(1), 1–26. http://www.jstor.org/stable/2958830

Ferrel, W. (1886). *Annual report of the Chief Signal Officer of the Army to the Secretary of War for the year 1886*. Government Printing Office.

Hanzer, F., Carmagnola, C. M., Ebner, P. P., Koch, F., Monti, F., Bavay, M., Bernhardt, M., François, H., Lehning, M., Morin, S., Olefs, M., & Strasser, U. (2020). Simulation of snow management in Alpine ski resorts using three different snow models. *Cold Regions Science and Technology*, *172*, 102995. https://doi.org/10.1016

Hedstrom, N. R., & Pomeroy, J. W. (1998). Measurements and modelling of snow interception in the boreal forest. *Hydrological Processes*, 12(10-11), 1611-1625 https://doi.org/10.1002

Lehning, M., Bartelt, P., Brown, B., Fierz, C., & Satyawali, P. (2002). A physical SNOWPACK model for the Swiss avalanche warning service. *Cold Regions Science and Technology*, 35(3), 123-145. https://doi.org/10.1016

Mann, H. B., & Whitney, D. R. (1947). On a test of whether one of two random variables is stochastically larger than the other. *Annals of Mathematical Statistics*, 18(1), 50-60. https://doi.org/10.1214

Murray, F. W. . (1967). On the Computation of Saturation Vapor Pressure. *Journal of Applied Meteorology and Climatology*, *6*(1), 203-204. https://doi.org/10.1175

Spandre, P., François, H., Verfaillie, D., Lafaysse, M., Deque, M., Eckert, N., George, E., & Morin, S. (2018). Climate Change Impacts On The Snow Reliability Of French Alps Ski Resorts. *Proceedings*, https://arc.lib.montana.edu/snow-science/item/2587

Storn, R., & Price, K. (1997). Differential evolution: A simple and efficient heuristic for global optimization over continuous spaces. *Journal of Global Optimization*, 11(4), 341-359. https://doi.org/10.1023

U.S. Government. (1976). *U.S. Standard Atmosphere, 1976*. National Oceanic and Atmospheric Administration, National Aeronautics and Space Administration, & United States Air Force.