

Weather-to-Key: Machine Learning for Planning and Operation of Satellite-QKD Ground Stations

Gianluca De Santis^{1,2}, Hannah Thiel^{3,4,5}, Lukas Tiefenthaler^{3,4,5}, Jasminder Sidhu⁶, Konstantin Kravtsov¹, Sana Amairi-Pyka¹, Eleni Diamanti², and James A. Grieve¹

¹Quantum Research Center, Technology Innovation Institute, PO Box 9639 Abu Dhabi, United Arab Emirates

²Sorbonne Université, CNRS, LIP6, 4 Place Jussieu, Paris F-75005, France

³Vigo Quantum Communication Center, University of Vigo, Vigo E-36310, Spain

⁴Escuela de Ingeniería de Telecomunicación, Department of Signal Theory and Communications, University of Vigo, Vigo E-36310, Spain

⁵AtlanTTic Research Center, University of Vigo, Vigo E-36310, Spain

⁶Dahlem Center for Complex Quantum Systems, Freie Universität Berlin, 14195 Berlin, Germany

Abstract

The deployment and operation of satellite quantum-key-distribution (QKD) networks require accurate knowledge of atmospheric turbulence, which determines optical coupling efficiency and therefore directly impacts the secure key rate. Conventional turbulence characterization is based on dedicated optical instrumentation and long-term monitoring campaigns, limiting both site assessment and operational forecasting. We investigate whether machine learning applied to routinely available meteorological observations can provide the turbulence information required to deploy and operate space-ground quantum networks. We address two complementary tasks. First, for ground-station siting, we predict the turbulence-regime distribution of new locations using only weather data. A model trained on a desert observatory reproduces the turbulence distribution of an unseen coastal site with 94% overlap and that of an unseen continental site with 88% overlap, substantially outperforming climatological baselines. Second, for operational scheduling, we develop an intra-night forecasting framework that predicts the evolution of the atmospheric turbulence throughout the remainder of the night. Combining complementary forecasting models yields 80% directional accuracy on the 69% of nights for which the models agree, providing a practical confidence measure for scheduling decisions. To connect these predictions to experimentally relevant quantum communication metrics, we map turbulence forecasts to secure key rates through an adaptive-optics-aware model of a satellite-QKD downlink. This enables candidate ground stations to be evaluated directly in terms of quantum-network performance. Our results demonstrate how machine learning can support both the planning and real-time operation of satellite-QKD infrastructure while expressing atmospheric predictions in the operational language of quantum information science.

1 Motivation

Satellite quantum key distribution has progressed from proof-of-principle demonstrations toward the deployment of operational quantum communication networks [1, 2]. As satellite-QKD systems move toward deployment, practical engineering questions become increasingly important. Network operators must decide where new optical ground stations should be located and, given multiple available stations, which one should be selected to receive a particular satellite pass. Both questions depend strongly on atmospheric turbulence. Atmospheric turbulence distorts the incoming optical wavefront, reducing single-mode-fibre coupling efficiency and consequently the photon detection rate, quantum bit error rate and secure key generation. To evaluate a possible site, dedicated instruments such as differential image motion monitors (DIMMs) must typically operate for months before a site's long-term characteristics are understood. Furthermore, historical measurements alone do not provide the forecasts needed for operational scheduling.

Existing atmospheric prediction models are typically heavily parameterised for specific locations and do not readily generalise across sites. Although weather observations do not directly measure optical turbulence, they capture the atmospheric conditions that drive its formation and evolution, making them attractive predictors of the Fried parameter. In this work, using weather and turbulence measurements collected at four observatories spanning three geographically distinct regions (Abu Dhabi, Vigo and Innsbruck), we investigate whether machine learning applied to ground-level weather observations can provide the turbulence information required by satellite-QKD systems. We consider two complementary prediction tasks centred on the same physical quantity: the Fried parameter r_0 , one of the standard measures of atmospheric turbulence strength. For infrastructure planning, we predict the long-term distribution of r_0 at previously unseen locations using weather observations alone. For network operation, we forecast the short-term evolution of r_0 throughout the night to support pass-scheduling decisions. Crucially, we connect both tasks to quantum communication performance through a turbulence-to-key-rate model incorporating adaptive optics. Unlike conventional turbulence-prediction studies, our objective is not atmospheric forecasting itself but its impact on secure quantum communication.

2 Predicting turbulence distributions for site selection

In the site-selection task, the learning target is the Fried parameter. Rather than predicting its instantaneous value, however, we seek the long-term distribution of r_0 , since this determines the probability that a station can sustain a quantum link and therefore influences both network design and adaptive-optics requirements. We formulate site selection as an out-of-distribution prediction problem. A gradient-boosted tree model [4] is first used for single site datasets to validate that the weather carries information about atmospheric turbulence. The same model is then trained on a reference site (Abu Dhabi) and transferred to geographically distinct sites (Vigo and Innsbruck), where no turbulence labels are available during training. The Fried parameter is discretised into five ordered turbulence regimes, and the model predicts the turbulence-regime distribution directly. A central challenge is preventing temporal leakage, which can dramatically inflate reported performance in atmospheric time-series data. Random train-test splits place strongly correlated observations from the same night on opposite sides of the split, producing overly optimistic estimates of generalisation. We therefore perform evaluation using complete-night partitions and remove features that implicitly encode site-specific calendar information. Hyperparameter optimization, calibration, and preprocessing are performed exclusively at the reference site prior to transfer. Using only weather observations, a model trained on a desert observatory reproduces the turbulence distribution of an unseen coastal site with 94% overlap, compared with 60% for a climatological baseline. Transfer to an unseen continental site achieves 88% overlap, compared with 57% for the baseline. These results indicate that weather-derived features capture predictive relationships that transfer across substantially different environments.

3 Intra-night forecasting for pass scheduling

For operational scheduling, the learning target remains the Fried parameter, but the objective shifts from long-term statistics to short-term forecasting. We therefore predict the evolution of r_0 throughout the remainder of the night from observations collected during its initial period. We combine two complementary predictive models. The first is a trajectory model based on a long short-term memory network that forecasts the evolution of r_0 at ten-minute resolution. The second is a random-forest classifier that predicts whether conditions will improve, degrade, or remain stable. Forecast evaluation again emphasises realistic deployment conditions. Nights affected by instrument failures are removed at the night level rather than by filtering individual samples, and forecasting targets are masked to accommodate gaps in observations without preferentially selecting calm nights. Evaluation is performed using grouped expanding-window splits across seasons. The resulting forecasts modestly

outperform persistence and seasonal climatology over short (≤ 3 h) horizons. Most importantly, agreement between the two models provides a useful confidence estimate. When both models predict the same trend, representing approximately 69% of nights, directional accuracy reaches 80%; when they disagree, performance falls substantially (a permutation test confirms the gap, $p < 0.001$). This agreement signal can therefore be used directly within scheduling policies, enabling operators to distinguish between high-confidence and low-confidence forecasts.

4 From turbulence prediction to quantum communication performance

The machine-learning models described above predict atmospheric quantities, whereas a quantum communication system ultimately cares about secure key generation. To bridge this gap, we construct a turbulence-to-key-rate transfer model for adaptive-optics-assisted satellite QKD. The model maps the Fried parameter to residual wavefront error, Strehl ratio, single-mode-fibre coupling efficiency [5], total optical transmission, and finally secure key rate for a decoy-state BB84 downlink [6]. Receiver aperture and adaptive-optics parameters are included explicitly, allowing different system architectures to be compared; the wavefront-error budget is calibrated against a published satellite-QKD adaptive-optics analysis [7]. This connection enables machine-learning predictions to be interpreted in operationally meaningful terms. Instead of asking how often turbulence exceeds an arbitrary threshold, one can ask how often a site supports a target secure key rate. Similarly, predicted turbulence distributions can be translated into adaptive-optics requirements for maintaining a desired level of quantum-link availability. For example, at an 0.8 m receiver aperture and a 1 kbit/s target key rate, the coastal site reaches 50% QKD-usable time only with six corrected Zernike radial orders, versus four for the desert site; a direct, physical quantification of the additional adaptive-optics investment the more turbulent location demands, obtained for the unseen site from weather alone. The machine-learning outputs therefore become directly relevant to engineering decisions regarding ground-station deployment and system design.

The contribution of this work is not atmospheric prediction in isolation, but the use of a common machine-learning framework centred on prediction the Fried parameter to support both long-term infrastructure planning and real-time operation of satellite quantum communication systems. By mapping weather observations to turbulence forecasts and ultimately to secure key generation, we obtain a weather-to-key pipeline for planning and operating future space-ground quantum networks. This illustrates how machine learning can transform readily available classical data into actionable engineering decisions for experimental quantum information science.

References

- [1] S.-K. Liao et al., “Satellite-to-ground quantum key distribution,” *Nature* **549**, 43–47 (2017).
- [2] J. Yin et al., “Entanglement-based secure quantum cryptography over 1,120 kilometres,” *Nature* **582**, 501–505 (2020).
- [3] D. L. Fried, “Optical resolution through a randomly inhomogeneous medium for very long and very short exposures,” *J. Opt. Soc. Am.* **56**, 1372–1379 (1966).
- [4] T. Chen and C. Guestrin, “XGBoost: A scalable tree boosting system,” in *Proc. 22nd ACM SIGKDD*, 785–794 (2016).
- [5] N. Jovanovic et al., “Efficient injection from large telescopes into single-mode fibres,” *Astron. Astrophys.* **602**, A57 (2017).
- [6] H.-K. Lo, X. Ma and K. Chen, “Decoy state quantum key distribution,” *Phys. Rev. Lett.* **94**, 230504 (2005).
- [7] V. Marulanda Acosta et al., “Analysis of satellite-to-ground quantum key distribution with adaptive optics,” *New J. Phys.* **26** (2024).