Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

When Does Fuzzy Adaptation Earn Its Complexity? A Calibration-Aware Computational Study of Finite-Shot Entanglement Detection

Som Subhro Nath

nathsomsubhro7@gmail.com

Independent Researcher https://orcid.org/0009-0007-6842-0019

Research Article

Keywords: Quantum entanglement detection, finite-shot measurements, fuzzy logic, adaptive measurement allocation, variance-based allocation, entanglement witnesses, detector calibration uncertainty, false-positive control, computational quantum physics

Posted Date: November 2nd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11541170/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

When Does Fuzzy Adaptation Earn Its Complexity? A Calibration-Aware Computational Study of Finite-Shot Entanglement Detection

Som Subhro Nath
Independent Researcher
Kolkata, West Bengal, India
E-mail ID: nathsomsubhro7@gmail.com
ORCID: 0009-0007-6842-0019

Abstract

Finite-shot entanglement detection requires measurement resources to be distributed across observable settings while balancing detection power, estimation error, and false-positive control. This study presents a comparative computational evaluation of three measurement-allocation strategies: fixed allocation, variance-adaptive allocation, and a fuzzy-adaptive strategy that combines weighted fuzzy membership scores with variance-based allocation. The methods are evaluated for a two-qubit Werner-type state family using a total budget of 3,000 measurement shots across four simulated detector conditions: ideal operation, fixed offset, uncertain calibration, and drift. The fuzzy strategy is selected from a prespecified family of 81 configurations using development-set detection power and evaluated on separate validation and test samples.

The results indicate that the relative performance of the allocation strategies depends on detector conditions and evaluation criteria. Detection power is high under ideal and fixed-offset conditions for all three methods, while uncertain calibration and drift are associated with reduced performance. Fuzzy adaptation achieves the highest observed detection power under uncertain calibration (98.59%) and drift (96.39%), but does not consistently minimize estimation error. Independent boundary-state validation also reveals imperfect false-positive-rate control: the fuzzy strategy exhibits a false-positive rate of 9.0% under uncertain calibration, exceeding the nominal 5% target. These findings demonstrate that apparent gains in detection power must be interpreted alongside false-positive behaviour and estimation accuracy. Because the analysis is computational and the configurations are selected using development-set performance, the results do not establish universal superiority or statistical significance of fuzzy adaptation. Instead, the study provides a condition-dependent benchmark of fuzzy-assisted measurement allocation and identifies calibration robustness and threshold control as important directions for further investigation.

Keywords: Quantum entanglement detection; finite-shot measurements; fuzzy logic; adaptive measurement allocation; variance-based allocation; entanglement witnesses; detector calibration uncertainty; false-positive control; computational quantum physics.

1. Introduction

Quantum entanglement is a fundamental resource in quantum information science, underpinning protocols in quantum communication, quantum computation, and quantum metrology. Yet the operational usefulness of an entangled state depends not merely on its theoretical existence, but on whether its entanglement can be established reliably from accessible experimental data. In realistic settings, this task is constrained by finite measurement budgets, statistical fluctuations, and imperfections in the measurement apparatus. These constraints make entanglement detection not only a problem in quantum theory, but also a problem of statistical estimation and resource allocation. The central question is therefore not simply whether a quantum state is entangled, but how reliably that property can be inferred when measurements are costly and the measurement process itself is imperfect.

Entanglement witnesses provide a practical means of addressing this problem without requiring complete quantum-state tomography. A witness is an observable whose expectation value satisfies a prescribed bound for all separable states, while a violation of that bound provides evidence of entanglement. For selected state families, witness-based detection can reduce the experimental requirements relative to full state reconstruction. The mathematical foundations of separability criteria and witness construction are well established [1–3]. Nevertheless, the practical reliability of witness-based inference depends on how accurately the relevant expectation values are estimated. Under finite sampling, empirical estimates fluctuate around their underlying values, and states close to the separability boundary are particularly challenging to classify. A decision rule that ignores these fluctuations can produce false-positive detections, while a conservative rule can fail to identify genuinely entangled states.

Measurement allocation is consequently an integral component of finite-shot entanglement detection. When a witness depends on several observables, the available measurement budget must be distributed among them. Uniform allocation is transparent and reproducible, but it does not explicitly account for differences in the statistical uncertainty associated with individual observables. Adaptive allocation attempts to use information obtained from preliminary measurements to determine how the remaining budget should be distributed. Related problems have been studied in quantum-state estimation and optimal experimental design, including adaptive procedures for finite-data estimation [4,5]. Such work establishes the relevance of measurement strategy to estimation efficiency, although its conclusions do not automatically extend to every witness-based detection problem.

The distinction between statistical uncertainty and systematic measurement error is equally important. Increasing the number of shots can reduce sampling variability, but it does not necessarily eliminate errors arising from calibration offsets, uncertain detector response, or temporal drift. Rosset et al. [6] showed that imperfect measurement settings can undermine quantum-state characterization and entanglement-witness evaluation. These considerations motivate the evaluation of measurement-allocation strategies beyond idealized sampling conditions. In particular, a method that performs well under one detector model may not retain that advantage when calibration uncertainty or drift changes the relationship between observed outcomes and the underlying quantum correlations.

Variance-adaptive allocation provides a natural statistical baseline for distributing a finite measurement budget. By using pilot estimates to approximate the relative uncertainty of different observables, such a procedure allocates shots according to an estimated uncertainty profile. Its appeal lies in its direct connection to statistical estimation and its relatively transparent allocation rule. However, pilot-based allocation remains dependent on the quality of preliminary estimates, the assumptions used to approximate uncertainty, and the relationship between the allocation objective and the final decision criterion. Reducing estimation uncertainty is not equivalent to controlling the false-positive rate of an entanglement test. The allocation rule and the threshold used

to classify a state must therefore be evaluated as components of a complete detection procedure rather than as independent guarantees of reliability.

This study investigates fuzzy logic as an alternative mechanism for mapping pilot-estimated uncertainty to measurement allocation. Fuzzy set theory represents graded membership in predefined categories rather than requiring every input to be assigned to a single crisp category [7]. In a fuzzy-adaptive allocation scheme, estimated uncertainty can contribute to categories such as low, medium, and high uncertainty to different degrees. Weighted membership scores can then be transformed into allocation proportions and combined with a variance-based allocation. The motivation for this approach is methodological rather than doctrinal: fuzzy logic offers a configurable mapping from estimated uncertainty to resource allocation, but its flexibility does not establish that it is statistically more efficient or more robust than a simpler rule. Its practical value must be determined by comparison against appropriate baselines under a clearly specified evaluation protocol.

A central issue is therefore whether fuzzy adaptation justifies the additional complexity it introduces. Membership functions, weighting parameters, and blending coefficients create design choices that may affect the resulting allocation. If configurations are selected using development-set performance, the selection procedure itself can influence the apparent performance of the chosen method. Furthermore, maximizing detection power alone may favour a configuration that also produces an unacceptably high false-positive rate. A meaningful comparison must consequently consider several outcomes, including detection power, false-positive behaviour, and estimation error. It must also distinguish performance on data used for configuration selection from performance evaluated on separate samples. Without these distinctions, an apparent improvement in detection can be difficult to interpret: it may reflect a genuine advantage in resource allocation, a change in decision-threshold behaviour, or sensitivity to the simulated conditions.

The present work addresses this question through a comparative computational study of finite-shot entanglement detection for a two-qubit Werner-type state family. We compare three allocation strategies: fixed allocation, variance-adaptive allocation, and a fuzzy-adaptive strategy that combines weighted fuzzy membership scores with variance-based allocation. Each strategy is evaluated under a common total budget of 3,000 measurement shots and across four simulated detector conditions: ideal operation, fixed offset, uncertain calibration, and drift. The analysis considers detection power, false-positive rates, and root-mean-square estimation error (RMSE). The fuzzy strategy is drawn from a prespecified family of 81 configurations, with configuration selection based on development-set detection power and subsequent evaluation on separate samples.

The contribution of this study is a condition-dependent comparison of a defined fuzzy-adaptive allocation procedure against fixed and variance-adaptive baselines within a common finite-shot detection framework. Rather than treating fuzzy logic as an intrinsically superior decision mechanism, we examine whether its use produces practically relevant differences across detector conditions and evaluation criteria. This framing also makes the limitations of the approach explicit: performance under simulated detector models does not establish robustness to all experimental imperfections, and high detection power does not by itself establish satisfactory false-positive control. The findings are therefore interpreted as comparative computational evidence, not as a general guarantee of superiority or a formal validation of fuzzy allocation for arbitrary quantum systems.

The remainder of the paper is organized as follows. Section 2 critically reviews the literature on entanglement witnesses, finite-shot detection, adaptive measurement allocation, detector imperfections, and fuzzy logic, and establishes the research gap addressed by this study. Section 3 presents the mathematical formulation and methodology, including the state family, entanglement witness, detector models, and allocation procedures. Section 4 describes the simulation design,

configuration-selection protocol, and evaluation metrics. Section 5 reports the comparative results across detector conditions. Section 6 discusses the findings, their implications, and the limitations of the proposed approach. Section 7 concludes the paper and identifies directions for future research.

2. Literature Review

2.1. Entanglement witnesses and finite-sample detection

Entanglement detection is conventionally formulated as a problem of distinguishing entangled quantum states from the set of separable states using experimentally accessible observables. Entanglement witnesses provide a particularly relevant framework because they permit entanglement to be certified without reconstructing the complete density operator. The theoretical foundations of separability criteria and witness construction were established by Horodecki et al. [1], Terhal [2], and subsequently reviewed in detail by Gühne and Tóth [3]. These contributions establish the mathematical basis of witness-based detection, but they do not, by themselves, resolve how finite measurement resources should be distributed when the relevant expectation values must be estimated experimentally.

For a Hermitian witness operator (W), a conventional detection criterion takes the form ($\mathrm{Tr}(W\rho) < 0$), where the witness is constructed so that its expectation value is nonnegative for all separable states in the relevant setting. Experimental evaluation replaces the exact expectation value with an estimate obtained from finite observations. The resulting decision is therefore statistical rather than purely algebraic. Even when the underlying state is separable, sampling fluctuations may yield a negative empirical witness value. Conversely, an entangled state may fail to be detected if the estimated witness does not cross the prescribed threshold. These two errors—false-positive detection and missed detection—are distinct and should not be conflated with estimation error.

The finite-data problem has been studied explicitly in quantum-state verification. Blume-Kohout et al. [8] investigated entanglement verification with finite data, emphasizing the role of statistical evidence when drawing conclusions from incomplete experimental samples. This line of work establishes that finite-sample entanglement claims require more than the evaluation of a point estimate. The uncertainty of the estimate and the rule used to convert it into a detection decision are integral to the inference procedure.

The distinction is consequential for measurement-allocation studies. An allocation strategy may reduce the mean squared error of a witness estimator without necessarily improving the operating characteristics of the resulting hypothesis test. Similarly, an increase in detection power is not sufficient evidence of methodological superiority if it is accompanied by an increase in the false-positive rate. A comparative analysis must therefore evaluate detection power, false-positive behaviour, and estimation error separately. The present study adopts this multi-criterion perspective rather than treating any single metric as a complete measure of detection quality.

2.2. Adaptive measurement and finite-shot allocation

Quantum-state estimation and experimental design provide the principal methodological background for allocating a finite number of measurements across observables. Nunn et al. [4] investigated optimal experiment design for quantum-state tomography, framing measurement selection in terms of fairness, precision, and experimental economy. Sugiyama et al. [5] studied adaptive experimental design for one-qubit state estimation with finite data, using a statistical update criterion related to average-variance optimality. Their numerical comparisons demonstrated that adaptive design can improve estimation precision relative to standard nonadaptive tomography under the conditions considered.

These studies establish that measurement design can be an active component of quantum estimation rather than a fixed experimental prerequisite. They do not, however, establish that every adaptive allocation rule improves every quantum-information task. Their objectives, measurement models, and state-estimation settings differ from the finite-shot witness-classification problem considered here. In particular, a design criterion that optimizes a state-estimation objective need not optimize the false-positive rate or detection power of a threshold-based entanglement test.

The allocation of shots among measurement terms has also been studied in variational quantum algorithms. Zhu et al. [9] proposed a variance-preserving shot-reduction strategy for variational quantum eigensolver measurements. Their work addresses the practical cost of estimating quantum expectation values and demonstrates the relevance of term-dependent shot assignment under a finite measurement budget. Liang et al. [10] investigated an artificial-intelligence-driven approach to shot reduction in variational quantum eigensolver calculations, using reinforcement learning to learn allocation policies. These studies show that shot allocation is an active methodological problem, with approaches ranging from statistical heuristics to learned policies.

The relevance of these results to the present work is methodological rather than direct. Variational quantum eigensolvers seek accurate objective-function estimates and successful optimization, whereas witness-based entanglement detection requires a decision concerning whether a state violates a separability criterion. The two problems share the constraint of finite measurement resources, but their loss functions and failure modes differ. An allocation strategy that reduces energy-estimation variance or accelerates variational convergence cannot be assumed to improve entanglement detection without a task-specific evaluation.

Variance-adaptive allocation serves as an important baseline in this context because it uses pilot observations to estimate the relative uncertainty of individual measurements. It is statistically motivated and comparatively transparent. Nevertheless, its performance depends on the reliability of pilot estimates and on the relationship between the estimated uncertainty and the final objective. The use of pilot measurements also creates a design dependency: preliminary data influence subsequent resource allocation, while the final estimator is based on a different, though related, sample. Any proposed alternative should therefore be compared against variance-adaptive allocation under a shared measurement budget and an explicitly stated evaluation protocol.

2.3. Imperfect measurements, calibration uncertainty, and drift

The assumption that measurement settings are implemented exactly is a significant limitation of idealized entanglement-detection models. Rosset et al. [6] examined the consequences of imperfect measurement settings for quantum-state tomography and entanglement witnesses. Their analysis showed that systematic errors arising from imperfect alignment of measurement bases can undermine the characterization of quantum resources and can affect witness-based conclusions. This establishes that measurement imperfections are not merely an additional source of sampling noise: they can alter the relationship between the intended observable and the measurement actually performed.

Morelli et al. [11] developed a framework for entanglement detection with imprecise measurements. They considered settings in which local measurements approximate, but do not exactly implement, their intended observables. Their analysis demonstrated that modest measurement inaccuracies can compromise established entanglement witnesses and derived corrections for a family of witnesses under bounded measurement inaccuracy. This work is directly relevant to the motivation for evaluating entanglement detection under imperfect detector conditions.

Cao et al. [12] extended this line of investigation to genuine multipartite entanglement detection, combining theoretical witness construction with an experimental demonstration under imperfect measurements. Their results further establish that measurement imperfections can materially affect entanglement certification and that explicit treatment of such imperfections can improve the robustness of witness-based inference. The setting differs from the two-qubit state family considered here, but the underlying methodological issue is shared: the validity of a detection conclusion depends on the relationship between the assumed measurement model and the physical measurement process.

These contributions should not be interpreted as interchangeable treatments of calibration offset, uncertain gain, and temporal drift. Imperfect measurement alignment, bounded inaccuracies in measurement settings, and time-varying detector response represent different error models. A method evaluated under one model is not thereby validated under the others. The present study consequently treats ideal operation, fixed offset, uncertain calibration, and drift as distinct simulated conditions rather than combining them into a single undifferentiated noise category.

A further distinction concerns what adaptive allocation can and cannot correct. Reallocating shots may reduce statistical uncertainty in selected observables, but it does not automatically remove systematic bias. If the measurement response is miscalibrated, additional shots can produce a more precise estimate of the wrong quantity. Calibration-aware evaluation must therefore examine not only estimation error but also the resulting classification behaviour. In this respect, detector imperfections and measurement allocation are related but separate methodological concerns: the former affect the measurement model, while the latter determine how finite resources are distributed within that model.

2.4. Fuzzy logic as a measurement-allocation mechanism

Fuzzy set theory, introduced by Zadeh [7], provides a mathematical framework for representing graded membership in predefined categories. Unlike crisp classification, in which an input belongs or does not belong to a category, fuzzy membership permits an input to have a degree of membership between zero and one. This representation is useful in decision systems where categories are deliberately defined to overlap or where a graded mapping is preferred to a rigid categorical boundary.

In a resource-allocation procedure, fuzzy membership functions can map estimated uncertainty to categories such as low, medium, and high. Weighted membership scores can then be normalized into allocation proportions. Such a procedure offers a configurable mapping from pilot-derived uncertainty to resource allocation. Its flexibility, however, is accompanied by additional design choices: the membership functions, weights, normalization procedure, and any blending with an alternative allocation policy must be specified. The resulting allocation is determined by these choices rather than by fuzzy logic in the abstract.

For the present study, the relevant question is not whether fuzzy systems can represent graded information; that capability follows from their mathematical construction. The relevant question is whether a particular fuzzy mapping of pilot-estimated uncertainty produces useful empirical differences relative to fixed and variance-adaptive allocation. This is a narrower and testable claim. The theoretical availability of fuzzy membership functions does not establish statistical efficiency, improved calibration robustness, or control of false-positive rates.

The distinction is especially important because the proposed strategy is a hybrid allocation procedure. It combines fuzzy membership scores with variance-based allocation through a blending parameter. Its behaviour therefore depends both on the fuzzy scoring mechanism and on the contribution of the variance-adaptive component. If the hybrid performs favourably in a given condition, the result cannot automatically be attributed to fuzzy logic alone. A stronger attribution

would require a design that isolates the effects of the fuzzy mapping, the variance-based component, and their combination.

The configuration-selection procedure introduces a further consideration. In the present study, the fuzzy family contains 81 configurations, and selection is based on detection power in development simulations. This protocol provides a systematic way to choose a configuration from a prespecified family, but it also means that the selected method is optimized for a particular criterion. Selection on detection power does not directly penalize elevated false-positive rates or higher estimation error. The evaluation must therefore report these outcomes independently and avoid interpreting the selected configuration as a universally optimal fuzzy controller.

2.5. Synthesis and research gap

The reviewed literature establishes several points relevant to the present investigation. First, entanglement witnesses provide a mathematically established route to detecting entanglement without complete state tomography [1–3]. Second, finite-sample estimation and adaptive experimental design demonstrate that measurement allocation can affect estimation performance under resource constraints [4,5,8]. Third, research on imperfect measurements establishes that systematic measurement inaccuracies can undermine witness-based inference, motivating evaluation under explicitly specified detector models [6,11,12]. Fourth, recent work on shot assignment demonstrates continuing interest in allocating measurement resources efficiently, although much of this work addresses variational quantum algorithms rather than threshold-based entanglement detection [9,10]. Finally, fuzzy set theory supplies a formal mechanism for constructing graded decision mappings, but it does not guarantee that a fuzzy allocation rule will outperform a statistical baseline [7].

These strands motivate a comparative question at their intersection: how does a prespecified fuzzy-adaptive allocation strategy perform relative to fixed and variance-adaptive allocation when a finite measurement budget is used for witness-based entanglement detection under different simulated detector conditions? The question is narrower than a claim to introduce a new theory of quantum measurement or to establish fuzzy logic as a superior allocation paradigm. It concerns the empirical behaviour of a defined allocation rule under a shared budget and a common evaluation framework.

The present study addresses this question through a computational comparison for a two-qubit Werner-type state family. It evaluates fixed allocation, variance-adaptive allocation, and a fuzzy-adaptive hybrid under ideal, fixed-offset, uncertain-calibration, and drift conditions. Detection power, false-positive rate, and root-mean-square estimation error are considered as distinct performance criteria. The fuzzy configuration is selected from a prespecified family using development-set detection power and then evaluated on separate samples.

The contribution should therefore be understood as a condition-dependent benchmark rather than a claim of universal methodological superiority. The reviewed literature supports the importance of finite-shot allocation, measurement imperfections, and explicit evaluation criteria. It does not, on the basis of the sources considered here, establish that the specific fuzzy-adaptive procedure evaluated in this study is superior to existing allocation methods. Nor does this review establish that no prior study has examined fuzzy allocation for a closely related quantum-measurement task. The defensible objective is to characterize the performance and trade-offs of the specified method, identify conditions in which its behaviour differs from the baselines, and assess whether the observed detection gains are accompanied by acceptable false-positive control and estimation accuracy.

3. Mathematical Formulation and Methodology

3.1. Quantum-state model

The study considers a two-qubit Werner-type state family defined by

$$\left[\rho(p) = p|\Phi^+\rangle\langle\Phi^+| + (1-p)\frac{I_4}{4},\right]$$

where

$$\left[|\Phi^+\rangle = \frac{|00\rangle + |11\rangle}{\sqrt{2}}, \qquad 0 \le p \le 1,\right]$$

and (I_4) denotes the identity operator on the two-qubit Hilbert space. The parameter (p) controls the contribution of the maximally entangled component relative to the maximally mixed state. For this family, the state is entangled if and only if $(p > 1/3)$.

The three Pauli-correlation observables considered are

$$[O_1 = X \otimes X, \qquad O_2 = Y \otimes Y, \qquad O_3 = Z \otimes Z.]$$

Their expectation values satisfy

$$\langle O_1\rangle_\rho = p,$$
$$\langle O_2\rangle_\rho = -p,$$
$$\langle O_3\rangle_\rho = p.$$

These observables provide the inputs to the entanglement-witness estimator. The analysis is restricted to this specified state family and measurement configuration; no claim is made that the allocation procedures are optimal for arbitrary two-qubit states or higher-dimensional systems.

3.2. Entanglement witness and decision rule

The study employs the witness

$$\left[W = \frac{1 - \langle XX\rangle + \langle YY\rangle - \langle ZZ\rangle}{4}.\right]$$

For the state family defined above,

$$\left[\mathbb{E}_{\rho(p)}[W] = \frac{1-3p}{4}.\right]$$

Consequently, $(W < 0)$ for $(p > 1/3)$, while the witness expectation is nonnegative for $(p \leq 1/3)$. The boundary state $(p = 1/3)$ is therefore the relevant separability-boundary case for evaluating false-positive behaviour.

Let $(\hat{c}_j)$ denote the estimated correlation associated with observable (O_j). The empirical witness is

$$\left[\widehat{W} = \frac{1 - \hat{c}_1 + \hat{c}_2 - \hat{c}_3}{4}.\right]$$

Detection is defined by the threshold rule

$$[D = \mathbb{1}\widehat{W} < \tau,]$$

where(τ) is a threshold estimated from simulated boundary-state data. Here, $(\mathbb{1}\cdot)$ denotes the indicator function. The threshold is not assumed to provide exact false-positive-rate control: it is estimated from a finite sample, and its realised operating characteristics must be assessed on separate validation data.

3.3. Finite-shot measurement allocation

The total measurement budget is fixed at

$$[N_{\text{total}} = 3000.]$$

The initial pilot stage allocates 30 shots to each of the three observables, giving

$$[N_{\text{pilot}} = 3(30) = 90.]$$

The remaining budget is therefore

$$[N_{\text{remaining}} = 2910.]$$

Let (n_j) denote the final number of shots assigned to observable (O_j). The allocation must satisfy

$$\sum_{j=1}^{3} n_j = 3000, \qquad n_j \geq 0,$$

with integer shot counts and the pilot observations included in the final total. The three methods differ in how they allocate the remaining budget after the pilot stage.

3.3.1. Fixed Allocation

The fixed method distributes the total budget approximately equally across the three observables. In the implemented configuration, this corresponds to 1,000 shots per observable. It serves as the nonadaptive baseline and does not use the pilot uncertainty estimates to change the allocation.

3.3.2. Variance-adaptive allocation

The variance-adaptive method estimates the relative uncertainty of each observable from the pilot data. Let $\left(\bar{y}_j^{(0)}\right)$denote the pilot sample mean and $\left(\hat{g}_j\right)$ the estimated detector gain. The pilot calculation uses

$$\left[\tilde{m}_j = \mathrm{clip}\left(\frac{\bar{y}_j^{(0)}}{\hat{g}_j}, -1,1\right),\right]$$

followed by the uncertainty proxy

$$\left[s_j = \sqrt{\frac{\max\left(1-\tilde{m}_j^2, 10^{-4}\right)}{n_j^{(0)}}},\right]$$

where $(n_j^{(0)} = 30)$ and $(\mathrm{clip}(\cdot))$ restricts its argument to the stated interval. The allocation proportions are then defined as

$$\left[a_j^{(V)} = \frac{s_j}{\sum_{k=1}^{3} s_k}.\right]$$

The remaining measurement budget is assigned according to these proportions, subject to the integer-allocation procedure implemented in the simulation.

This rule is a specified uncertainty-based heuristic. Its use does not establish that the resulting allocation minimizes the variance of the final witness estimator under every detector model.

3.3.3. Fuzzy-adaptive allocation

The fuzzy-adaptive method transforms normalized pilot uncertainty into weighted membership scores. First, define

$$\left[u_j = \frac{s_j}{\max\limits_k s_k}.\right]$$

The low-, medium-, and high-uncertainty membership functions are

$$[\mu_L(u) = \mathrm{clip}(1-2u, 0,1)\,,]$$

$$[\mu_M(u) = \mathrm{clip}(1-|2u-1|, 0,1)\,,]$$

$$[\mu_H(u) = \mathrm{clip}(2u-1, 0,1)\,.]$$

The weight vector is selected from the prespecified set

$$[\boldsymbol{w} = (w_L, w_M, w_H), \qquad w_r \in 1,2,4.]$$

For each observable, the weighted fuzzy score is

$$[q_j = w_L\mu_L(u_j) + w_M\mu_M(u_j) + w_H\mu_H(u_j).]$$

To maintain positive scores before normalization, define

$$[\tilde{q}_j = \max(q_j, 10^{-8}).]$$

The fuzzy allocation proportions are

$$\left[a_j^{(F)} = \frac{\tilde{q}_j}{\sum_{k=1}^{3}\tilde{q}_k}.\right]$$

Finally, the fuzzy proportions are combined with the variance-adaptive proportions:

$$\left[a_j^{(FA)} = (1-\lambda)a_j^{(F)} + \lambda a_j^{(V)},\right]$$

where

$$[\lambda \in 0.25,0.50,0.75.]$$

The resulting proportions are nonnegative and sum to one. This follows directly from the normalization of the component allocations and the constraint ($0 \le \lambda \le 1$). It is a feasibility property of the construction, not a claim of improved statistical performance.

The configuration family contains ($3^3 = 27$) weight vectors and three blending coefficients, giving

$$[|\mathcal{F}| = 3^3 \times 3 = 81]$$

candidate configurations. The fuzzy method is therefore not an unconstrained search over all possible fuzzy systems; it is a comparison of a prespecified finite family.

3.4. Detector conditions and calibration correction

Four simulated detector conditions are considered: ideal operation, fixed offset, uncertain calibration, and drift. These conditions are used to examine whether the relative performance of the allocation methods changes when the measurement process departs from the idealized setting.

The final correlation estimates apply the detector gain and offset corrections to the combined measurement data. Conceptually, the correction has the form

$$\left[\hat{c}_j = \frac{\bar{y}_j - \hat{o}_j}{\hat{g}_j},\right]$$

where $(\bar{y}_j)$ is the combined observed sample mean, $(\hat{o}_j)$ is the estimated offset, and $(\hat{g}_j)$ is the estimated gain. The corrected correlations are then used to calculate $(\widehat{W})$.

An implementation detail requires explicit qualification. The pilot uncertainty calculation divides the pilot mean by the estimated gain and clips the result, but uses a zero offset in that intermediate

calculation. The final estimator, in contrast, applies the estimated offset as well as the gain correction. Thus, the pilot uncertainty calculation and final estimator do not apply identical calibration corrections. The possible effect of this discrepancy on allocation and subsequent performance has not been isolated by the current experiment. The results must therefore be interpreted as the performance of the implemented procedure, including this inconsistency, rather than as validation of a fully offset-corrected pilot strategy.

3.5. Configuration selection and evaluation separation

The fuzzy configuration is selected from the 81 candidates using detection power on the development simulations. The selection criterion is

$$\frac{1}{N_{\mathrm{dev}}}\sum_{i=1}^{N_{\mathrm{dev}}} \mathbb{1}\,\widehat{W}_i(\boldsymbol{\theta}) < \tau,$$

Where $(\boldsymbol{\theta})$ represents a candidate weight vector and blending coefficient. The selected configuration maximizes this development-set criterion. When multiple configurations attain the same maximum, the implemented configuration order determines which candidate is retained.

The threshold is estimated using a separate simulated boundary-state sample, with 250 threshold samples specified by the experimental design. The selected allocation method is subsequently evaluated using separate validation and test samples. The nominal false-positive target is

$$[\alpha_{\mathrm{FPR}} = 0.05.]$$

The false-positive rate is estimated as the proportion of separable-state trials classified as entangled. Independent boundary-state validation is reported separately from the broader separable-state test evaluation, since these datasets address related but distinct aspects of false-positive behaviour.

Selection based only on development-set detection power has a direct interpretive consequence: it does not optimize false-positive rate or estimation error. Those quantities are therefore treated as independent evaluation criteria rather than as properties guaranteed by the selection procedure.

3.6. Performance measures

The methods are compared using three principal measures.

Detection power:

For entangled-state test trials, detection power is the fraction classified as entangled:

$$\left[\hat{\pi} = \frac{1}{N_E}\sum_{i=1}^{N_E} D_i\,,\right]$$

where (N_E) is the number of entangled-state trials.

False-positive rate:

For separable-state trials, the false-positive rate is

$$\left[\widehat{\mathrm{FPR}} = \frac{1}{N_S}\sum_{i=1}^{N_S} D_i\,,\right]$$

where (N_S) is the number of separable-state trials. This metric evaluates how often the decision rule incorrectly identifies an entangled state when the tested state is separable.

Root-mean-square error

Estimation accuracy is assessed using

$$\left[\mathrm{RMSE} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}\left(\widehat{W}_i - W_i\right)^2}\,,\right]$$

where (W_i) is the theoretical witness expectation for trial (i), and $(\widehat{W}_i)$ is its empirical estimate.

These measures are reported separately because they quantify different properties. High detection power does not imply low false-positive rates, and low RMSE does not by itself guarantee correct threshold-based classification.

3.7. Scope of the methodological claims

The methodology defines a finite, prespecified comparison among three allocation procedures under a fixed measurement budget and four simulated detector conditions. Its purpose is to characterize empirical performance differences, not to establish a universal optimum.

No formal optimality claim is made for the fuzzy allocation. Nor does the normalization of its allocation proportions establish robustness to calibration error. Any claim of practical advantage must be supported by the evaluation results and interpreted alongside the configuration-selection procedure, the finite validation samples, and the pilot offset inconsistency described above.

4. Experimental Design and Evaluation Protocol

4.1. Experimental objectives and comparison structure

The experimental design evaluates whether fuzzy-adaptive allocation produces a measurable difference in finite-shot entanglement detection relative to fixed and variance-adaptive allocation. The comparison is conducted under a common measurement budget and across four simulated detector conditions. The objective is not to establish universal optimality, but to characterize the dependence of detection and estimation performance on allocation strategy and detector model.

Three methods are evaluated:

1. **Fixed allocation**, which distributes the measurement budget equally across the three observables.

2. **Variance-adaptive allocation**, which uses pilot-derived uncertainty estimates to determine allocation proportions.
3. **Fuzzy-adaptive allocation**, which combines a weighted fuzzy membership score with the variance-adaptive allocation through a prespecified blending coefficient.

The fuzzy method is selected from a finite family of 81 configurations. The same selected configuration is used for the corresponding detector condition during its subsequent evaluation. The comparison therefore concerns a development-selected fuzzy policy, not an oracle configuration chosen retrospectively from test performance.

4.2. Simulation parameters and state sampling

The simulation uses a total measurement budget of 3,000 shots per trial. Each trial begins with a pilot stage of 30 shots for each of the three observables, consuming 90 shots. The remaining 2,910 shots are assigned according to the relevant allocation strategy. The measurement budget is held constant across methods to prevent differences in total resource consumption from confounding the comparison.

The entangled-state test set uses six values of the state parameter,

$$[p \in 0.4,0.5,0.6,0.7,0.8,0.9.]$$

The separable-state test set uses

$$p \in \left\{0,\ 0.1,\ 0.2,\ \frac{1}{3}\right\}.$$

These values sample the two sides of the witness boundary and include the boundary state ($p = 1/3$). The resulting evaluation is restricted to the specified Werner-type family and the selected parameter values. It should not be interpreted as a representative sample of all entangled and separable two-qubit states.

The reproducibility script specifies 500 entangled test trials and 300 separable test trials per method and detector condition. The reported summary contains 498 entangled observations, reflecting the realized evaluation dataset. This discrepancy between the nominal test count and the exported count must be reconciled against the trial-level files before the final manuscript is released. No unverified explanation for the difference is assumed here.

4.3. Detector conditions

Four detector conditions are evaluated:

- **Ideal**: a reference condition without the modelled detector imperfections.
- **Fixed offset**: a reference condition in which the detector response includes a fixed offset.
- **Uncertain calibration**: a condition in which calibration parameters are uncertain.
- **Drift**: a condition in which detector response changes according to the specified drift model.

The conditions are analysed separately. This is necessary because fixed offset, calibration uncertainty, and drift represent different deviations from the ideal measurement model and need not affect estimation or classification in the same way.

The ideal condition provides a baseline for evaluating the allocation procedures in the absence of modelled detector imperfections. The fixed-offset condition isolates a different form of measurement imperfection, while uncertain calibration and drift probe the methods under more complex detector

assumptions. These simulations do not establish robustness to arbitrary experimental noise, unmodeled correlations, or time-dependent measurement effects beyond those encoded in the implementation.

The precise detector-generation equations and calibration-estimation procedures are part of the executable experimental specification. They should be reproduced directly from the authoritative script in the final version rather than replaced with generic noise equations.

4.4. Fuzzy-configuration search

The fuzzy family is prespecified by three weights,

$$[w_L, w_M, w_H \in 1,2,4,]$$

and a blending coefficient,

$$[\lambda \in 0.25,0.50,0.75.]$$

The full candidate set therefore contains 81 configurations. Configuration selection is performed separately for each detector condition using development-set detection power. The development design specifies 180 entangled trials per condition.

For a candidate configuration $(\boldsymbol{\theta})$, the development criterion is

$$\frac{1}{N_{\mathrm{dev}}} \sum_{i=1}^{N_{\mathrm{dev}}} \mathbb{1}\, \widehat{W}_i(\boldsymbol{\theta}) < \tau.$$

The selected configuration is the candidate with the largest observed development-set detection power. Ties are resolved according to the order in which configurations appear in the prespecified search family. This deterministic rule makes the selection procedure reproducible, but it does not make tied configurations statistically distinguishable.

The selection criterion is deliberately narrow. It does not jointly optimize detection power, false-positive rate, and RMSE. Consequently, the subsequent test and validation results must be interpreted as an assessment of a power-selected configuration across multiple performance criteria, not as the outcome of multi-objective optimization.

4.5. Threshold calibration and independent evaluation

The detection threshold is estimated from simulated data at the separability boundary. The implementation specifies

$$[N_{\mathrm{threshold}} = 250]$$

threshold-calibration samples and a nominal false-positive target of

$$[\alpha_{\mathrm{FPR}} = 0.05.]$$

For each method and detector condition, the threshold is obtained from the lower (α_{FPR}) quantile of the simulated boundary-state witness estimates. A trial is classified as entangled when its estimated witness is strictly below this threshold.

The use of a separate boundary-state sample for threshold estimation is intended to separate threshold calibration from subsequent evaluation. The realized false-positive rate is nevertheless a random quantity because the threshold is estimated from a finite sample. Accordingly, the nominal 5% target is a calibration objective, not a guarantee that the realized false-positive rate will equal 5%.

The evaluation includes two distinct assessments of false-positive behaviour. First, the general test set includes separable states at several values of (p), including $(p = 1/3)$. Second, an independent boundary-state validation experiment uses 600 trials per method and detector condition. The latter specifically assesses the false-positive behaviour at the separability boundary. Results from these two evaluations are reported separately because their state compositions and statistical purposes differ.

4.6. Evaluation metrics and uncertainty estimates

Three principal metrics are used.

Detection power is the proportion of entangled test trials classified as entangled. It measures sensitivity to the entangled states represented in the test design.

False-positive rate is the proportion of separable trials classified as entangled. It is reported both for the separable test set and for the independent boundary-state validation set. The two estimates are not interchangeable: the former aggregates over the selected separable-state distribution, whereas the latter focuses on the boundary state.

Root-mean-square error is calculated between the estimated witness and its theoretical expectation. Entangled-state and separable-state RMSE are reported separately, allowing estimation accuracy to be examined across the two test populations.

The implementation computes percentile bootstrap intervals for selected performance estimates. Such intervals describe sampling variability under the resampling procedure used; they do not, by themselves, account for every source of uncertainty introduced by searching across 81 fuzzy configurations. The selection process must therefore remain explicit when interpreting the reported intervals.

The comparisons are consequently interpreted descriptively. Differences in point estimates are not presented as statistically significant effects, and overlapping intervals are not interpreted as proof of equivalence.

4.7. Reproducibility and implementation

The experiment is implemented in Python using NumPy, pandas, and Matplotlib. The reproducibility script specifies a base seed of 20261025 and uses NumPy SeedSequence to construct deterministic random streams from explicit integer identifiers. This avoids dependence on Python's process-dependent hash behaviour for random-seed generation. The script also records software and platform metadata.

The exported outputs include summary-level performance metrics, configuration-selection records, boundary-state validation results, trial-level evaluation data, and configuration metadata. These artifacts are intended to support independent inspection of the allocation procedures and reported outcomes.

4.8. Interpretive constraints

The experimental protocol supports a comparative evaluation of three specified allocation strategies under a finite measurement budget and four simulated detector conditions. It does not establish that

fuzzy logic is generally superior to variance-based allocation, nor that the selected fuzzy configuration is optimal outside the evaluated search family.

Three constraints are especially relevant to interpretation. First, the fuzzy configurations are selected using development-set detection power, which may favour high sensitivity without guaranteeing false-positive control. Second, the threshold is estimated from a finite simulated boundary-state sample, so the realized false-positive rate can depart from its nominal target. Third, the pilot uncertainty calculation and final estimator do not apply identical offset corrections. The possible impact of this mismatch has not been separately quantified.

These constraints do not invalidate the comparison, but they define what the experiment can support. The results should be read as evidence about the implemented allocation procedures under the specified simulation assumptions, not as a general guarantee of detection reliability under arbitrary detector imperfections.

5. Results

5.1. Overview of the comparative evaluation

The three allocation procedures were evaluated under four detector conditions: Ideal, FixedOffset, UncertainCalibration, and Drift. Each trial used a nominal budget of 3,000 shots. The exported summary contains 498 entangled-state trials and 300 separable-state trials for each method–condition combination. The entangled-state count is two below the nominal test count of 500 specified in the experimental protocol. This discrepancy is reported explicitly and remains to be reconciled against the implementation.

Performance was assessed using detection power on entangled states, false-positive rate on separable states, and root mean squared error (RMSE). These metrics characterize different aspects of performance and must not be treated as interchangeable. Higher detection power does not establish better calibration, and lower RMSE does not establish valid false-positive control.

Across the four conditions, detection power was high for all three procedures. The ranking of methods nevertheless varied across detector conditions and evaluation criteria. The results therefore do not support a general claim that fuzzy adaptation is superior to fixed or variance-adaptive allocation.

5.2. Detection power

The held-out detection-power estimates are reported in *Table 1*. Values are percentages; the intervals in parentheses are the reported percentile confidence intervals.

Table 1. Held-out detection power for entangled states.

Detector condition	Fixed	Variance-adaptive	Selected fuzzy-adaptive
Ideal	99.80 (99.40–100.00)	99.80 (99.40–100.00)	99.60 (99.00–100.00)
FixedOffset	99.80 (99.40–100.00)	99.60 (99.00–100.00)	99.80 (99.40–100.00)
UncertainCalibration	97.79 (96.39–99.00)	97.99 (96.78–99.20)	98.59 (97.39–99.60)
Drift	96.18 (94.38–97.79)	95.98 (94.18–97.59)	96.39 (94.58–97.99)

Under Ideal detection, Fixed and VarianceAdaptive both achieved 99.80% power, while the selected fuzzy procedure achieved 99.60%. Under FixedOffset, Fixed and fuzzy adaptation both achieved 99.80%, marginally exceeding the 99.60% observed for VarianceAdaptive. These differences are small and do not establish a meaningful advantage for any one procedure.

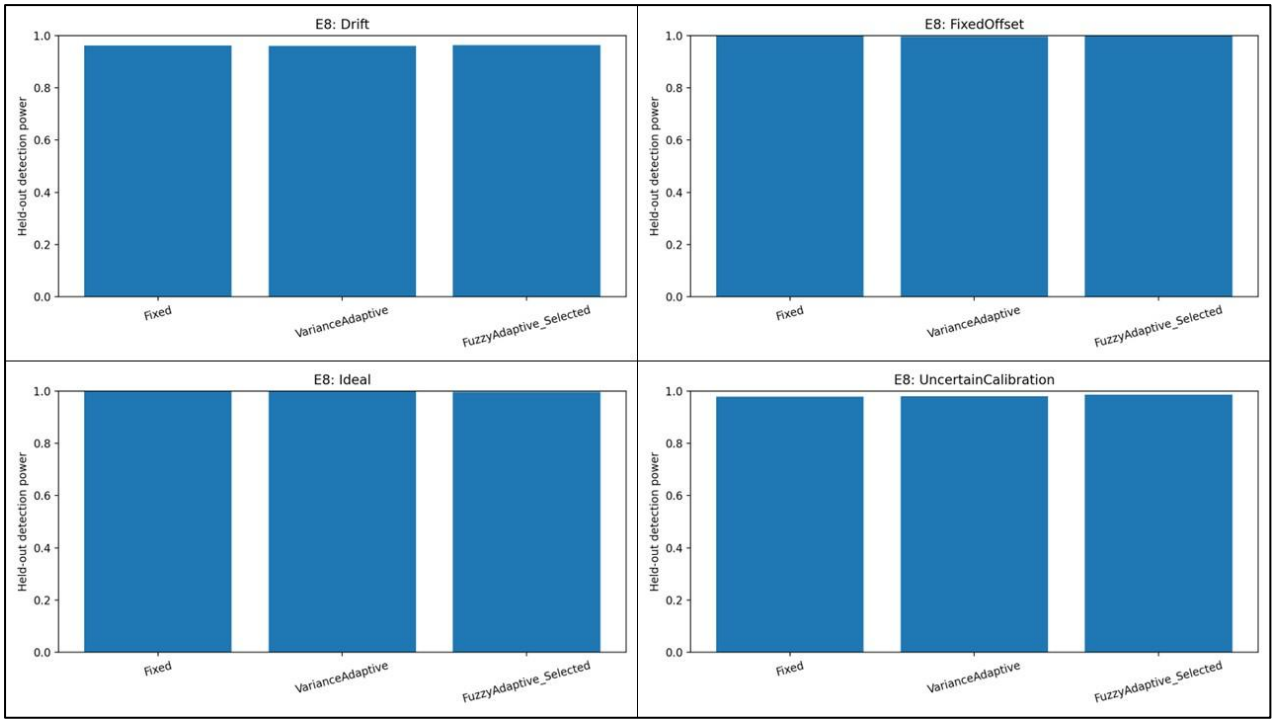


Fig.1. Detection power across detector conditions

The largest observed power advantage for fuzzy adaptation occurred under UncertainCalibration. Its power was 98.59%, compared with 97.79% for Fixed and 97.99% for VarianceAdaptive. The corresponding absolute differences were 0.80 and 0.60 percentage points. Under Drift, fuzzy adaptation again produced the highest observed power, but the differences were only 0.21 percentage points relative to Fixed and 0.41 percentage points relative to VarianceAdaptive.

All three procedures showed lower power under UncertainCalibration and Drift than under Ideal detection. Fuzzy adaptation did not remove the effects of detector imperfection. Furthermore, because the fuzzy configuration was selected using development-set detection power, the observed advantage cannot establish that the fuzzy rule is intrinsically optimal or that it would retain the same ranking under other state families or detector models.

5.3. False-positive rates on separable states

False-positive rates on the held-out separable-state distribution are reported in *Table 2*. Values are percentages; the intervals in parentheses are the reported percentile confidence intervals.

Table 2. Held-out false-positive rates on separable states.

Detector condition	**Fixed**	**Variance-adaptive**	**Selected fuzzy-adaptive**
Ideal	1.00 (0.00–2.33)	1.00 (0.00–2.33)	3.00 (1.33–5.00)
FixedOffset	2.33 (0.67–4.33)	1.00 (0.00–2.33)	1.00 (0.00–2.33)
UncertainCalibration	2.00 (0.67–3.67)	2.00 (0.67–3.67)	1.33 (0.33–2.67)
Drift	1.33 (0.33–2.67)	2.33 (0.67–4.33)	2.00 (0.67–3.67)

The separable-state results do not yield a consistent ranking. Under Ideal detection, fuzzy adaptation produced a higher observed false-positive rate (3.00%) than Fixed and VarianceAdaptive (both 1.00%). Under FixedOffset, Fixed had the highest rate (2.33%), while both adaptive procedures produced 1.00%. Under UncertainCalibration, fuzzy adaptation had the lowest observed rate (1.33%),

compared with 2.00% for each alternative. Under Drift, Fixed had the lowest rate (1.33%), followed by fuzzy adaptation (2.00%) and VarianceAdaptive (2.33%).

The configuration-selection criterion is relevant to interpreting these results. Fuzzy configurations were selected on development-set detection power, not through a joint objective that penalized false positives. The lower observed false-positive rate under UncertainCalibration is therefore a favourable empirical result, but it does not show that the selection procedure optimized false-positive control. Conversely, the higher rate under Ideal detection demonstrates that the selected fuzzy configuration did not consistently improve specificity.

Each method–condition combination was evaluated on 300 separable-state trials. Consequently, the false-positive estimates are based on a limited number of events. They should be interpreted as descriptive comparisons within the simulated setting, not as precise guarantees of error rates in deployment.

5.4. Independent boundary-state null validation

The separable-state test results and the boundary-state null-validation results address different questions. The former summarizes false positives over the specified separable-state test distribution; the latter evaluates performance at the entanglement boundary using 600 validation trials per method and condition. Independent boundary-state null-validation results are reported in *Table 3*. The two sets of estimates are not pooled.

*Table 3. Independent boundary-state null-validation false-positive rates***.**

Detector condition	Fixed	Variance-adaptive	Selected fuzzy-adaptive
Ideal	4.33 (2.83–6.00)	3.50 (2.17–5.00)	5.83 (4.00–7.67)
FixedOffset	6.33 (4.50–8.33)	4.33 (2.83–6.00)	4.83 (3.33–6.67)
UncertainCalibration	7.33 (5.33–9.50)	6.50 (4.50–8.50)	9.00 (6.67–11.33)
Drift	5.50 (3.67–7.50)	5.67 (3.83–7.67)	7.17 (5.17–9.33)

Note: Values are percentages. Parenthetical intervals are the reported percentile confidence intervals. The nominal false-positive reference level is 5%.

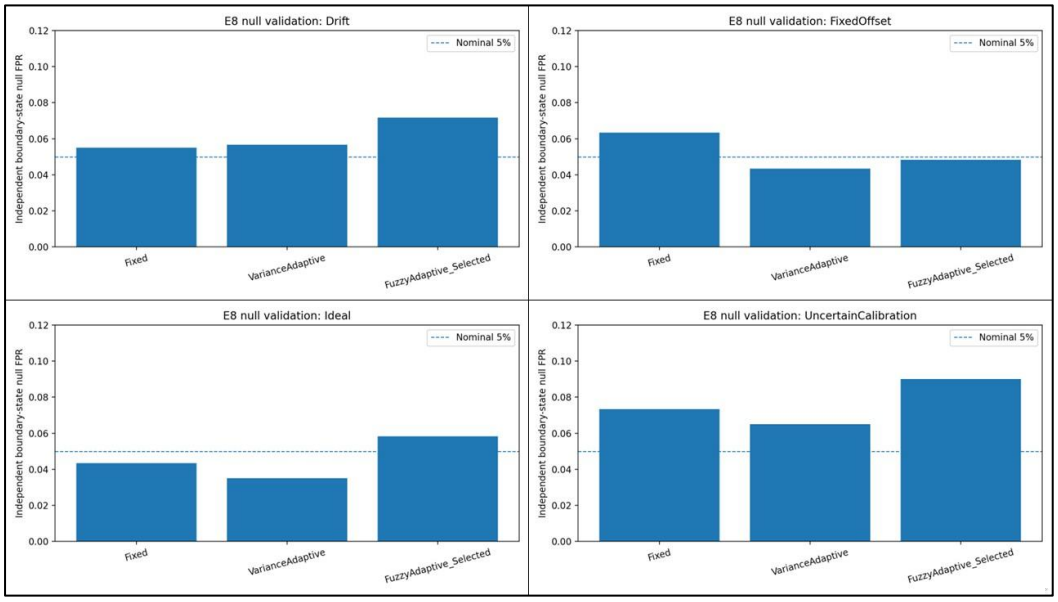


Fig. 3. Independent boundary-state false-positive validation

The boundary-state results are less favourable than the aggregate separable-state test results. Under Ideal detection, fuzzy adaptation produced a null false-positive rate of 5.83%, compared with 4.33% for Fixed and 3.50% for VarianceAdaptive. Under FixedOffset, the fuzzy procedure yielded 4.83%, below Fixed at 6.33% but above VarianceAdaptive at 4.33%.

The largest adverse result occurred under UncertainCalibration. Fuzzy adaptation produced a boundary-state false-positive rate of 9.00%, compared with 7.33% for Fixed and 6.50% for VarianceAdaptive. All three observed rates exceeded the nominal 5% reference. The reported interval for fuzzy adaptation was 6.67–11.33%, entirely above that reference. Within this simulation, the selected fuzzy configuration did not maintain the intended nominal false-positive level under UncertainCalibration.

Under Drift, fuzzy adaptation again had the highest observed boundary-state false-positive rate: 7.17%, compared with 5.50% for Fixed and 5.67% for VarianceAdaptive. The reported intervals for all three procedures include 5%; therefore, these results do not establish that each underlying false-positive probability differs from the nominal level. Nevertheless, the higher point estimate for fuzzy adaptation is not evidence of improved calibration.

These results constrain the interpretation of the power comparison. Fuzzy adaptation achieved the highest observed power under UncertainCalibration, but it also produced the highest boundary-state false-positive rate in that condition. This is consistent with a trade-off that cannot be evaluated from detection power alone. Because the configuration-selection criterion did not explicitly constrain boundary-state false positives, the results do not support a claim that fuzzy adaptation improves error control.

5.5. Estimation error

Entangled-state estimation errors, measured by RMSE, are reported in *Table 4*. Lower values indicate smaller error under the reported metric.

Table 4. RMSE on entangled-state trials.

Detector condition	Fixed	Variance-adaptive	Selected fuzzy-adaptive
Ideal	0.01090	0.01202	0.01013
FixedOffset	0.02323	0.02625	0.02223
UncertainCalibration	0.02594	0.02864	0.02669
Drift	0.01919	0.02123	0.01910

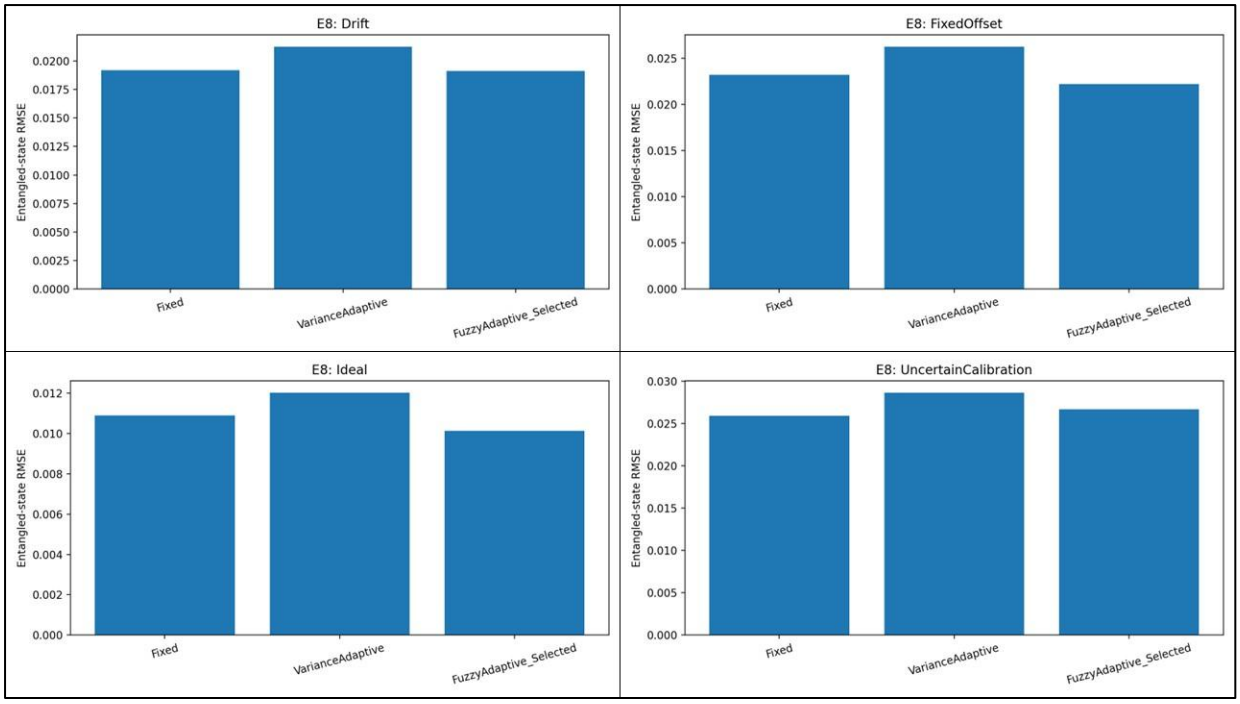


Fig.2. Entangled-state estimation error across detector conditions

Under Ideal detection, fuzzy adaptation achieved the lowest entangled-state RMSE (0.01013), followed by Fixed (0.01090) and VarianceAdaptive (0.01202). Under FixedOffset, fuzzy adaptation again had the lowest RMSE (0.02223), compared with 0.02323 for Fixed and 0.02625 for VarianceAdaptive.

The ranking changed under UncertainCalibration. Fixed achieved the lowest RMSE (0.02594); fuzzy adaptation produced 0.02669, and VarianceAdaptive had the largest error (0.02864). Under Drift, fuzzy adaptation and Fixed were close, with RMSE values of 0.01910 and 0.01919, respectively. VarianceAdaptive again produced the largest error (0.02123).

The RMSE results provide limited evidence that fuzzy adaptation can improve estimation accuracy under some simulated conditions. They do not support a general claim of lower error under detector imperfection. Notably, fuzzy adaptation did not minimize RMSE under UncertainCalibration, the condition in which it achieved its largest observed detection-power advantage.

RMSE is also conditional on the state family, detector model, estimator, and sampling procedure used in this experiment. It does not, by itself, establish unbiasedness, calibration validity, or generalization to other quantum states. The reported values therefore do not justify treating the fuzzy allocation rule as a universally more statistically efficient estimator.

5.6. Selected fuzzy configurations

The selected configuration varied across detector conditions, as summarized in *Table 5*.

*Table 5. Fuzzy configurations selected using development-set detection power***.**

Detector condition	Membership weights	Variance blend
Ideal	(1, 1, 1)	0.25
FixedOffset	(1, 1, 1)	0.25

UncertainCalibration	(1, 4, 1)	0.75
Drift	(4, 4, 4)	0.75

The search covered the prespecified family of 81 fuzzy configurations. Selection was based on development-set detection power, with ties resolved by the prescribed configuration order. The resulting configurations were condition-dependent: Ideal and FixedOffset selected equal weights with a smaller variance blend, whereas UncertainCalibration and Drift selected a larger blend. The membership weights also differed between the latter conditions.

This variation indicates that the search did not identify a single configuration that dominated across all detector conditions. Instead, it selected different candidates under different simulated environments. Under Ideal and FixedOffset conditions, multiple candidates attained the same development-set power; deterministic tie-breaking therefore influenced the final selection. The selected parameters should not be interpreted as uniquely identified or physically optimal.

The UncertainCalibration configuration, with weights (1, 4, 1) and blend 0.75, achieved the highest observed held-out power among the compared procedures in that condition. It also produced the highest boundary-state null false-positive rate in that condition. These outcomes are compatible because the selection criterion rewarded detection power without imposing an explicit boundary-state false-positive constraint. The result illustrates the limitations of power-only configuration selection when error control is a primary requirement.

5.7. Summary of empirical findings

Four conclusions are supported by the reported simulations.

First, all three procedures achieved high detection power on the evaluated entangled-state test distribution. Second, the selected fuzzy procedure had the highest observed power under UncertainCalibration and Drift, but the absolute advantages were modest. Third, RMSE rankings depended on detector condition; fuzzy adaptation did not consistently minimize estimation error. Fourth, independent boundary-state validation revealed elevated false-positive rates for the selected fuzzy procedure under UncertainCalibration and Drift, with the UncertainCalibration point estimate reaching 9.00%.

Taken together, these findings do not establish general superiority of fuzzy adaptation. They support a narrower conclusion: within this finite-shot simulation, a prespecified fuzzy allocation family can yield competitive detection power, but selection by power alone does not guarantee lower estimation error or nominal boundary-state false-positive control. Stronger claims would require a selection objective that explicitly incorporates error control, further independent validation, and evaluation across additional state families and detector models.

6. Discussion

6.1. Fuzzy-adaptive allocation as a condition-dependent strategy

The present study evaluates whether fuzzy adaptation can improve finite-shot entanglement detection relative to fixed and variance-adaptive allocation under explicitly modeled detector conditions. The results establish that the selected fuzzy procedure is competitive across the evaluated settings, with its relative performance depending on the detector regime and the evaluation metric. They do not establish universal superiority, nor do they support dismissing fuzzy adaptation on the basis of any single unfavourable result.

The clearest power advantage occurred under UncertainCalibration. The selected fuzzy procedure achieved 98.59% detection power, compared with 97.79% for Fixed and 97.99% for VarianceAdaptive. Under Drift, fuzzy adaptation achieved 96.39%, compared with 96.18% and 95.98%, respectively. These results show that the proposed allocation mechanism can retain high detection power in the presence of the modeled detector imperfections and can outperform the two baselines on this metric in selected conditions.

The magnitude of these differences must nevertheless be kept in perspective. The observed power advantage under Drift was small, and the Ideal and FixedOffset conditions did not show a consistent advantage for fuzzy adaptation. The evidence therefore supports a condition-dependent performance claim rather than a claim of general dominance.

The central contribution is the comparative evaluation itself: fuzzy allocation is assessed not merely by whether it can attain high detection power, but by how its performance changes across detector conditions and across distinct statistical criteria. This broader evaluation exposes behaviour that a power-only comparison would fail to capture.

6.2. What the fuzzy allocation mechanism contributes

The fuzzy procedure combines pilot-derived uncertainty information, membership-based weighting, and a blend with variance-adaptive allocation. Its configuration is selected from a prespecified family rather than being fixed in advance as one universal rule. This design provides a structured way to investigate whether a fuzzy representation of allocation priorities can yield competitive measurement-budget distributions.

The selected configurations varied across detector conditions. Under Ideal and FixedOffset, the selected weights were (1, 1, 1), with a variance blend of 0.25. Under UncertainCalibration, the selected weights were (1, 4, 1), with a blend of 0.75. Under Drift, the selected weights were (4, 4, 4), also with a blend of 0.75.

This variation is consistent with the condition-dependent behaviour of the procedure. The search did not identify a single configuration that was selected in every environment. Instead, different configurations were retained under different simulated detector conditions, reflecting differences in their development-set performance.

However, this observation must be interpreted precisely. The results demonstrate the operation of the specified configuration-selection procedure; they do not establish that the selected weights have a unique physical interpretation or that the configuration family contains the globally optimal allocation rule. Under Ideal and FixedOffset, multiple candidates achieved the same development-set power, meaning that deterministic tie-breaking influenced the final selection.

The study therefore establishes the feasibility of the proposed fuzzy allocation approach within the specified simulation framework. Whether its membership structure contributes benefits beyond variance-based allocation and blending remains a distinct question. The reported performance belongs to the complete procedure, including its configuration search, rather than to fuzzy membership functions considered in isolation.

6.3. Detection power under detector imperfections

The detection-power results show that all three procedures perform strongly on the evaluated entangled-state test distribution. Under Ideal and FixedOffset conditions, their observed power was approximately 99.6–99.8%. Under UncertainCalibration and Drift, power declined for all three methods.

Fuzzy adaptation produced its largest observed advantage under UncertainCalibration, where its power exceeded that of Fixed by 0.80 percentage points and that of VarianceAdaptive by 0.60 percentage points. Under Drift, the corresponding differences were 0.21 and 0.41 percentage points. These results indicate that the proposed procedure can be competitive in conditions where detector imperfections affect performance, although the size of the observed advantage varies substantially.

The high power achieved by all methods under Ideal detection also imposes a ceiling on the improvements that can be observed. In such conditions, small numerical differences are not compelling evidence of meaningful practical superiority. The more challenging detector conditions are consequently informative, but they do not remove the need to consider estimation error and false-positive behaviour.

Importantly, the deterioration under UncertainCalibration and Drift was not unique to any one allocation procedure. The results indicate that detector conditions remain material determinants of detection performance. Fuzzy adaptation did not eliminate their effects; it retained competitive performance within them.

6.4. Estimation accuracy and false-positive control

Detection power alone provides an incomplete assessment of an entanglement-detection procedure. The RMSE and null-validation results demonstrate why.

The fuzzy procedure achieved the lowest entangled-state RMSE under Ideal and FixedOffset conditions. Under Drift, its RMSE was marginally lower than that of Fixed. Under UncertainCalibration, however, Fixed achieved the lowest RMSE. These results show that the allocation rule's relative estimation performance depends on the detector condition. They do not support a general claim that fuzzy adaptation minimizes estimation error.

The independent boundary-state validation provides a further qualification. Under UncertainCalibration, the selected fuzzy procedure achieved a null false-positive rate of 9.00%, compared with 7.33% for Fixed and 6.50% for VarianceAdaptive. Under Drift, the corresponding fuzzy rate was 7.17%, again higher than the two alternatives. Thus, the conditions in which fuzzy adaptation retained a modest power advantage also included conditions in which its boundary-state false-positive performance was less favourable.

This result is central to the interpretation of the proposed method. The fuzzy configuration was selected on development-set detection power, not on an objective that jointly optimized power and boundary-state error control. The higher observed power under UncertainCalibration is therefore consistent with the selection criterion, while the elevated null false-positive rate reveals a dimension of performance that the selection criterion did not directly optimize.

These outcomes do not establish that fuzzy allocation inherently produces higher false-positive rates. They establish that the particular fuzzy configurations selected by the present procedure did not consistently achieve the best boundary-state error control. Equally, the favourable RMSE results in some conditions do not establish that fuzzy adaptation corrects detector bias or provides calibrated uncertainty estimates.

The appropriate interpretation is therefore multidimensional: the proposed method achieves competitive detection power, while its estimation accuracy and boundary-state false-positive behaviour remain condition-dependent.

6.5. Significance for finite-shot measurement allocation

The study contributes to the evaluation of finite-shot entanglement detection by examining allocation as part of a complete statistical decision procedure. The measurement budget is finite, and the allocation rule affects how observations are distributed across the selected observables. Yet the final decision also depends on detector correction and threshold calibration.

The results show that a competitive allocation rule cannot be judged solely by the number of entangled states it detects. A method may retain high power while producing less favourable boundary-state false-positive behaviour. Similarly, a lower RMSE under one detector condition does not guarantee better performance under another.

The distinction between aggregate separable-state false positives and boundary-state null validation is particularly important. The separable-state test results showed comparatively low false-positive rates across the methods. The boundary-state evaluation, however, exposed elevated rates under some detector conditions, particularly for the selected fuzzy procedure under UncertainCalibration. These results are not contradictory: the two evaluations characterize different parts of the decision problem.

Accordingly, the present work demonstrates the value of evaluating finite-shot allocation procedures across complementary criteria. It also illustrates why adaptation should be assessed together with the estimator, detector model, and decision threshold, rather than treated as an isolated optimization problem.

The scope of this conclusion remains specific to the simulated Werner-type two-qubit state family, the observables used, the 3,000-shot budget, the four detector conditions, and the specified configuration-selection procedure. The results do not establish performance for arbitrary quantum states, multipartite systems, alternative entanglement witnesses, or physical hardware whose behaviour differs from the modelled detector conditions.

6.6. Limitations

Some limitations should be stated explicitly to delimit the conclusions of the study.

First, the exported summary reports 498 entangled-state test observations, whereas the experimental protocol specifies a nominal target of 500. The reported results describe the exported data. The discrepancy must be reconciled against the implementation before the manuscript is finalized; the nominal count should not be substituted for the observed count.

Second, configuration selection was based on development-set detection power. Although the final comparison used held-out results, the selection criterion did not explicitly constrain boundary-state false-positive rates. The independent validation results demonstrate the practical relevance of this distinction. The present findings should therefore be interpreted as the performance of a power-selected fuzzy procedure, not as the outcome of a procedure designed to guarantee a prescribed false-positive level.

Third, the boundary-state validation used 600 trials per method and detector condition. The resulting false-positive estimates are subject to Monte Carlo uncertainty. The reported percentile intervals are descriptive and do not, by themselves, establish formal superiority or equivalence between methods. The study did not use a confirmatory hypothesis-testing framework with multiplicity adjustment; numerical rankings should not be presented as statistically established method-level differences.

Fourth, the decision threshold was estimated from a finite set of simulated boundary-state samples. The realized false-positive rate consequently depends on the calibration procedure and the modelled detector conditions. A nominal target of 5% does not guarantee a realized rate of 5% in every condition, as the validation results illustrate.

Fifth, the pilot allocation calculation and final estimator handle offset information differently. The pilot calculation uses a zero-offset assumption, whereas the final estimator uses an estimated offset. The quantitative effect of this inconsistency has not been isolated. The results should therefore not be interpreted as evidence that the allocation rule alone explains all observed differences.

Finally, the study is computational and model-specific. Its conclusions are conditional on the state family, detector models, shot budget, estimator, and random-sampling protocol used. The results do not demonstrate that the assumed detector models encompass all relevant imperfections in physical quantum experiments, nor do they establish generalization beyond the evaluated setting.

These limitations qualify the scope of the findings; they do not negate the comparative evidence obtained within the stated experimental design.

6.7. Broader Interpretation

The findings position fuzzy-adaptive allocation as a viable candidate for finite-shot entanglement detection rather than as a replacement for established allocation procedures. Its strongest observed advantage was in detection power under UncertainCalibration, with a smaller advantage under Drift. Its estimation-error performance was favourable in some conditions but not uniformly so, and its boundary-state false-positive performance was less favourable under UncertainCalibration and Drift.

The results also clarify the role of configuration selection. A fuzzy family can adapt its allocation priorities through condition-specific parameter selection, but the quality of the resulting procedure depends on the objective used to select those parameters. When selection rewards detection power alone, the selected configuration need not provide the best balance between detection and false-positive control.

The contribution of this study is therefore not a claim that fuzzy logic universally outperforms variance-based allocation. It is a structured computational comparison showing that fuzzy-adaptive allocation can achieve competitive power, while also revealing the trade-offs that emerge when estimation error and boundary-state false positives are examined independently.

This is a substantive result for evaluating adaptive measurement allocation. It identifies both the circumstances in which the proposed procedure performs competitively and the criteria under which that performance must be qualified.

6.8. Concluding Discussion

The study demonstrates that the specified fuzzy-adaptive procedure can attain high detection power under a finite measurement budget and remain competitive with fixed and variance-adaptive baselines across four simulated detector conditions. Its largest observed power advantage occurred under UncertainCalibration, while the RMSE and boundary-state validation results showed that this advantage did not extend uniformly to every performance criterion.

The evidence supports a precise conclusion: fuzzy adaptation is a promising allocation strategy within the evaluated simulation framework, but its performance is condition-dependent and its superiority is not universal. In particular, competitive detection power should not be conflated with lower estimation error or reliable nominal false-positive control.

By evaluating these outcomes separately, the present work provides a more complete characterization of fuzzy-adaptive allocation than a comparison based on detection power alone. The resulting contribution is empirical and methodological: it establishes the performance profile of a specified fuzzy allocation procedure, identifies the trade-offs associated with power-based configuration selection, and delineates the conditions under which stronger claims would exceed the available evidence.

7. Conclusion

This study investigated whether fuzzy-adaptive measurement allocation provides a measurable advantage over fixed and variance-adaptive allocation for finite-shot entanglement detection under detector imperfections. Using a Werner-type two-qubit state family, a fixed budget of 3,000 shots, four simulated detector conditions, and a prespecified family of 81 fuzzy configurations, the study compared detection power, estimation error, and false-positive behaviour.

The results establish that the evaluated fuzzy-adaptive procedure is competitive, but not uniformly superior. Its highest observed detection power under UncertainCalibration was 98.59%, compared with 97.79% for fixed allocation and 97.99% for variance-adaptive allocation. Under Drift, the corresponding power was 96.39%. These improvements were modest and did not translate into consistent advantages across the remaining metrics.

The estimation-error results were similarly condition-dependent. Fuzzy adaptation achieved the lowest entangled-state RMSE under Ideal and FixedOffset conditions and a marginally lower RMSE than fixed allocation under Drift. Under UncertainCalibration, fixed allocation produced the lowest RMSE. More importantly, independent boundary-state validation revealed a 9.00% false-positive rate for the selected fuzzy procedure under UncertainCalibration, exceeding the nominal 5% reference and the corresponding observed rates of both baseline methods. Under Drift, fuzzy adaptation also produced the highest observed boundary-state false-positive rate among the three procedures.

These findings expose the central methodological result of the study: competitive detection power does not, by itself, establish the superiority of an adaptive measurement-allocation procedure. The behaviour of the complete detection pipeline depends on the allocation rule, configuration-selection criterion, estimator, detector model, and threshold-calibration procedure. In the present implementation, selection based exclusively on development-set detection power did not guarantee favourable estimation error or nominal boundary-state false-positive control.

The contribution is therefore a comparative computational characterization of a specified fuzzy-adaptive allocation procedure, rather than a claim of universal superiority for fuzzy systems. The experiments demonstrate that fuzzy allocation can achieve high detection power within the evaluated simulation framework, while identifying conditions under which its performance must be qualified. The conclusions remain restricted to the tested state family, detector models, shot budget, and selection protocol; generalization to other quantum systems or physical hardware is not established.

Within these bounds, the study provides evidence that fuzzy-adaptive allocation merits consideration as a candidate strategy for finite-shot entanglement detection. Its value must be judged by the joint behaviour of detection power, estimation error, and false-positive control—not by the strongest performance metric in isolation.

8. References

1. M. Horodecki, P. Horodecki, and R. Horodecki, “Separability of mixed states: Necessary and sufficient conditions,” *Physics Letters A*, vol. 223, no. 1–2, pp. 1–8, 1996. https://doi.org/10.1016/S0375-9601(96)00706-2
2. B. M. Terhal, “Bell inequalities and the separability criterion,” *Physics Letters A*, vol. 271, no. 5–6, pp. 319–326, 2000. https://doi.org/10.1016/S0375-9601(00)00401-1
3. O. Gühne and G. Tóth, “Entanglement detection,” *Physics Reports*, vol. 474, no. 1–6, pp. 1–75, 2009. https://doi.org/10.1016/j.physrep.2009.02.004
4. J. Nunn, B. J. Smith, G. Puentes, I. A. Walmsley, and J. S. Lundeen, “Optimal experiment design for quantum state tomography: Fair, precise, and minimal tomography,” *Physical Review A*, vol. 81, 042109, 2010. https://doi.org/10.1103/PhysRevA.81.042109
5. T. Sugiyama, P. S. Turner, and M. Murao, “Adaptive experimental design for one-qubit state estimation with finite data based on a statistical update criterion,” *Physical Review A*, vol. 85, 052107, 2012. https://doi.org/10.1103/PhysRevA.85.052107
6. D. Rosset, R. Ferretti-Schöbitz, J.-D. Bancal, N. Gisin, and Y.-C. Liang, “Imperfect measurement settings: Implications for quantum state tomography and entanglement witnesses,” *Physical Review A*, vol. 86, 062325, 2012. https://doi.org/10.1103/PhysRevA.86.062325
7. L. A. Zadeh, “Fuzzy sets,” *Information and Control*, vol. 8, no. 3, pp. 338–353, 1965. https://doi.org/10.1016/S0019-9958(65)90241-X
8. R. Blume-Kohout, J. O. S. Yin, and S. J. van Enk, “Entanglement verification with finite data,” *Physical Review Letters*, vol. 105, 170501, 2010. https://doi.org/10.1103/PhysRevLett.105.170501
9. L. Zhu, S. Liang, C. Yang, and X. Li, “Optimizing shot assignment in variational quantum eigensolver measurement,” *Journal of Chemical Theory and Computation*, vol. 20, no. 6, pp. 2390–2403, 2024. https://doi.org/10.1021/acs.jctc.3c01113
10. S. Liang, L. Zhu, X. Liu, C. Yang, and X. Li, “Artificial-intelligence-driven shot reduction in quantum measurement,” 2024. https://doi.org/10.1063/5.0219663
11. S. Morelli, H. Yamasaki, M. Huber, and A. Tavakoli, “Entanglement detection with imprecise measurements,” *Physical Review Letters*, vol. 128, 250501, 2022. https://doi.org/10.1103/PhysRevLett.128.250501
12. H. Cao, S. Morelli, L. A. Rozema, C. Zhang, A. Tavakoli, and P. Walther, “Genuine multipartite entanglement detection with imperfect measurements: Concept and experiment,” *Physical Review Letters*, vol. 133, 150201, 2024. https://doi.org/10.1103/PhysRevLett.133.150201

9. Data and Code Availability

The datasets, source code, configuration files, and graphical outputs associated with this study are publicly available through Zenodo at https://doi.org/10.5281/zenodo.23736634. The repository provides the computational materials supporting the reported analyses and is intended to facilitate independent inspection and reproducibility of the results.

10. Statements and Declarations

10.1 Conflict of Interests:

No conflict of interests were identified in relation to the content, methodology, or findings of this study. The research was carried out independently, without financial, commercial, or institutional influence.

10.2 Funding Information:

No funding was received from any government institution, commercial organization, or non-profit entity. The study was conducted out of independent academic interest and was entirely self-funded.

10.3 Research Involving Humans and/or Animals:

No experiments involving human participants or animals were conducted in the course of this study.

10.4 Informed Consent:

Informed consent was not applicable, as no human participants were involved in the study.

10.5 Author Contributions:

The author conceived and designed the study, developed the experimental framework, conducted the experiments and analysis, and wrote and revised the manuscript.

11. Acknowledgements

Sincere acknowledgement is extended to **Mr. Ashim Kumar Nath, Bose Institute**, whose constant encouragement and intellectual influence have been instrumental in sustaining a rigorous engagement with scientific inquiry. His emphasis on critical questioning, methodological discipline, and evidence-based reasoning has provided enduring motivation throughout the conception and development of this work