




Adaptive quantum sensing for wave-like dark matter searches

arXiv: [2609.xxxxx]
In collaboration with



Codex

Wave dark matter signals at direct detection

❖ Coherent oscillation of classical wave dark matter (DM)

$$a(t) \simeq a_0 \sin(m_{\text{DM}}t - m_{\text{DM}}\vec{v}_{\text{DM}} \cdot \vec{x} + \delta)$$

❖ Unknown properties of DM signals

- Frequency = m_{DM}
- Amplitude $\propto a_0$
- Oscillation phase = δ
- Direction

ex) Dark photon polarization


Axion-induced $\vec{B}_{\text{eff}}$

Example: “Detection of hidden photon dark matter using the direct excitation of transmon qubits”
[Chen, et al. PRL '23]

❖ Resonant drive

as example of quantum sensing

Hajime's talk (Tue.)
Marianne's talk (Fri.)



$$m_{\text{DM}} = \omega_{\text{qubit}}$$

Population readout and projective measurements

❖ Population readout

- “DM signal flips the qubit spin at probability p ”


❖ Quantum description

$$|\psi\rangle = \cos\frac{\theta}{2}|g\rangle + \sin\frac{\theta}{2}e^{i\phi}|e\rangle$$

- $p = \sin^2\frac{\theta}{2}$
- Additional information from ϕ

❖ Projective measurement

- DOF of readout-axis choice
- “population” = z -projective measurement



Population vs transverse readout

❖ Population readout


- 👍 Insensitive to unknown phase ϕ
- 👎 Quadratic response $2\sin^2(\theta/2) \propto \theta^2$

❖ Transverse oracle (ϕ -readout)

- 👍 Linear response $\propto \theta$
- 👎 Need to know ϕ

❖ ϕ is determined by signal direction

- For DM signal direction $\tilde{\phi}$
 $\phi = \tilde{\phi} - \pi/2$ is unknown $\Rightarrow$ learn it!



Adaptive strategy

❖ During repeated measurements with n total shots,

Population readout for initial steps

- Phase-insensitive signal search

Phase learning


- Insert trial shots (**information gain**) to learn the signal phase φ

Converge to transverse oracle

- The optimal **ϕ -readout** for signal detection

❖ Help reject false positives/look-elsewhere effects

Bayesian adaptive readout loop



Bayesian update and final decision

❖ Bayesian update

- Start with $q_0 = \frac{1}{2}$ and $\pi_0(\theta, \tilde{\varphi})$
- If measurement along $\vec{n}_k$ results in $r_k = \pm 1$

$$q_{k+1} = \frac{q_k m_k(r_k | \vec{n}_k)}{(1 - q_k) p_{0,k}(r_k | \vec{n}_k) + q_k m_k(r_k | \vec{n}_k)}$$

Similar for $\pi_k(\theta, \tilde{\varphi})$

The diagram shows the Bayesian update equation. The denominator consists of two terms: $(1 - q_k) p_{0,k}(r_k | \vec{n}_k)$ and $q_k m_k(r_k | \vec{n}_k)$. Arrows point from labels below to these terms: 'When signal absent' points to $(1 - q_k)$, 'Prob. of r_k ' points to $p_{0,k}(r_k | \vec{n}_k)$, and 'When signal exists' points to q_k . Another 'Prob. of r_k ' label points to $m_k(r_k | \vec{n}_k)$. The text 'Similar for $\pi_k(\theta, \tilde{\varphi})$ ' is to the right of the equation.

❖ Final decision: frequentist test based on

$$\log B_{10,n} \equiv \log \frac{q_n}{1 - q_n}$$

Utilities 1: information gain

❖ This is exploration utility to learn signal existence and parameters

❖ Use “mutual information”



- $I(\Lambda_k; R_k)$

= How much readout $R_k \in \{\pm 1\}$ tells us about **signal existence** & **parameters** Λ_k

= $I(\text{signal existence}; R_k) + q_k I(\theta, \tilde{\varphi}; R_k)$



Inclines to parameter learning @ $q_k \rightarrow 1$



Utilities 2: Helstrom

❖ This is detection utility to collect signal evidence


❖ Helstrom measurement:

- Discrimination of ρ_0 and ρ_1
- Optimal measurement is
$$\vec{n}_{\text{opt}} \propto \vec{r}_0 - \vec{r}_1$$
- Converges to ϕ -readout for weak signal




Challenges due to unknown ϕ

$$\text{e.g. } \vec{r}_0 - \int \frac{d\phi}{2\pi} \vec{r}_1(\phi) \propto \vec{z}$$




Detector imperfections enter through likelihood



Two benchmark setups


1. Noisy baseline setup

- $\mathcal{C} = 0.99$
- $\epsilon = 0.005$
- Fake excitation rate by noise: $p_{z,0} \simeq 1\%$



2. Higher-fidelity setup



- $\mathcal{C} = 0.999$
- $\epsilon = 0.0005$
- Fake excitation rate by noise: $p_{z,0} \simeq 0.1\%$



Question: How detector qualities change the optimal learning-detection balance?



Adaptive trajectory (higher-fidelity)

- ❖ Clear environment allows efficient localization of posterior
- ❖ Hybrid of Information gain/Helstrom is empirically the best




Adaptive trajectory (noisy baseline)

- ❖ Noisy environment makes signal learning more valuable
- ❖ Pure information gain is empirically the best




Visualizing look-elsewhere effects (noisy baseline)

❖ Adaptive trajectory for noise-only model



Fake signal

Fake signal



Statistical treatment

❖ Composite hypothesis test


- H_0 : noise only
- H_1 : noise + signal ($\theta, \tilde{\varphi}$)
- Signal marginalized by $\pi_0(\theta, \tilde{\varphi})$

❖ Asymmetric testing

- Type-I error (false-positive) $\alpha < 5\%$
- Maximize signal detection power

❖ Run pseudo-experiments to

- Determine threshold from $\alpha = 5\%$
- Evaluate detection power



Non-adaptive baselines

❖ z-readout



Phase insensitive



Quadratic response

❖ x-readout



Linear response



Has a blind spot $\hat{n}_{\text{sig}} \simeq \hat{x}$

❖ Alternative x/y-readout




Linear response



Lose ½ statistics


Also, transverse oracle with ϕ -readout



Phase-averaged detection power (higher-fidelity)

❖ Best strategy depends on n


1. z-readout for small n
2. Adaptive for large n



Phase-averaged detection power (higher-fidelity)

❖ Best strategy depends on n

1. z -readout for small n
2. Adaptive for large n




Transverse wins z-readout (higher-fidelity)

❖ Best strategy is

1. z-readout for $n < n_{z:\text{adapt}}$
2. Adaptive for $n > n_{z:\text{adapt}}$

❖ Crossings due to different scaling

P	$n_{z:O\perp}$	$n_{z:x}$	$n_{z:xy}$	$n_{z:\text{adapt}}$
0.50	2.97×10^2	3.20×10^3	3.99×10^3	1.53×10^3
0.70	5.17×10^2	9.50×10^3	5.54×10^3	1.99×10^3
0.90	9.41×10^2	3.21×10^5	8.22×10^3	3.12×10^3



Transverse wins z-readout (noisy baseline)


❖ Best strategy is

1. z-readout for $n < n_{z:\text{adapt}}$
2. Adaptive for $n > n_{z:\text{adapt}}$

❖ Crossings due to different scaling

P	$n_{z:O_1}$	$n_{z:x}$	$n_{z:xy}$	$n_{z:\text{adapt}}$
0.50	7.88	2.32×10^2	3.03×10^2	1.47×10^2
0.70	7.90	7.20×10^2	3.77×10^2	2.12×10^2
0.90	7.50	3.08×10^4	4.89×10^2	2.86×10^2

❖ Adaptive sensing more important for noisy environments



Some information theory of asymmetric testing

❖ Stein's lemma tells type-II error (false negative rate) $\beta \rightarrow \exp[-\xi n + o(n)]$

- For non-adaptive approaches



Setting	Exponent ξ	ξ for small signals
Classical	$D_{\text{KL}}(P_0 P_1)$	$\propto$ Fisher information F_C
Quantum with qubit resets	$D_{\text{meas}}(\rho_0 \rho_1)$	$\propto$ Fisher information F_Q
Quantum with collective meas.	$D(\rho_0 \rho_1)$	$\propto$ Fisher information F_{BKM}

❖ Adaptive advantage proven for subset of composite problems [Bergh, et al. '23]

- Problems with non-convex set of classical channels
- Problems with non-convex set of quantum channels (convex is open)

An evidence for adaptive advantage (noisy baseline)

❖ Adaptive strategy gives the best Stein exponent



Experimental implications and next steps

❖ Implication to decision making


- Adaptive is best for $n \gtrsim n_{z:\text{adapt}}$
- $n_{z:\text{adapt}}$ is determined by **detector properties**

❖ Online update, NOT offline analysis

- Info discarded by measurement never recovered

❖ Next steps

- Finite signal coherence time
- Global optimization $\Rightarrow$ **DRL**
- Optimize initial state, interrogation time, etc. $\Rightarrow$ **VQA**



“Realizing a deep reinforcement learning agent discovering real-time feedback control strategies for a quantum system”

[Reuer, et al. Nat. Comm. '23]


Conclusion

❖ Main message

- Adaptive strategy can be beneficial for DM searches
- The best search strategy determined by detector setups; deserves consideration during experimental design

❖ Fascinating interdisciplinary direction

Particle × Quantum × ML



Backup slides

Imperfection channels

❖ Let $\rho = \frac{1}{2}(1_2 + \vec{r} \cdot \vec{\sigma})$ be the ideal quantum state

1. Relaxation

- $\mathcal{R}_T(\rho) = \frac{1}{2} [1_2 + \eta_2(T)(r_x\sigma_x + r_y\sigma_y) + \{1 - \eta_1(T)(1 - r_z)\}\sigma_z]$
- For $T_1 = T_2 = T$, $\eta_2(T) \simeq 0.63$, $\eta_1(T) \simeq 0.53$

2. Finite contrast


- $\mathcal{D}_C(\rho) = C\rho + \frac{(1-C)}{2} 1_2 \quad (0 \leq C \leq 1)$

3. Readout error rate

- $\mathcal{B}_\epsilon(r|r') = (1 - \epsilon)\delta_{r,r'} + \epsilon\delta_{r,-r'} \quad \left(0 \leq \epsilon \leq \frac{1}{2}\right)$


Detuning effects

❖ Detuning Δ makes optimal readout time-dependent: $\phi(t) = \varphi + \pi/2 + \Delta t$



Repeated myopic optimization $\neq$ global optimization


❖ Fixed- z is the best with unknown phase $\Leftrightarrow$ No phase information is learned



Phase-averaged detection power (noisy baseline)

❖ Best strategy for increasing n_{shots}

1. z-readout for small n_{shots}
2. Adaptive for large n_{shots}



Geometric polarization diagnostic

❖ Phase information efficiency

- $f_{\text{pol}} \equiv 2 \left\langle \left\{ \hat{n}_{\text{meas}} \cdot (\hat{z} \times \hat{n}_{\text{sig}}) \right\}^2 \right\rangle$
$$= \begin{cases} 0 & \text{(fixed-z)} \\ 1 & \text{(fixed-x/y)} \\ 2 & \text{(oracle)} \end{cases}$$
- Parameterizes how efficiently phase information was used

