

Hierarchical disorder in moiré exciton photoluminescence probed by spectral-descriptor correlations

Katsunori Wakabayashi*

Research Center for Materials Nanoarchitectonics (MANA), *National Institute for Materials Science (NIMS)*,
 Namiki 1-1, Tsukuba 305-0044, Japan



(Received 1 June 2026; accepted 17 July 2026; published 7 August 2026)

Hyperspectral photoluminescence maps of moiré transition-metal dichalcogenide heterobilayers encode rich information about the underlying disorder, but extracting it is hampered by the ambiguity of assigning individual spectral peaks. We show that the spatial correlations among a set of simple, peak-decomposition-free spectral descriptors—such as the centroid energy, the dominant-peak energy, and a sharp-line fraction—provide a peak-decomposition-free route to infer features of that disorder. Different descriptors act as filters that select different components of a multiscale disorder landscape—a smooth, micron-correlated background and a dense set of localized traps—and therefore acquire different spatial correlation lengths. The central prediction is a correlation-length hierarchy, $\xi(E_{\text{cent}}) \geq \xi(E_{\text{dom}})$, which we derive by splitting the dominant-peak energy into a smooth background part and a short-range trap-switching fluctuation. The same picture explains the measured interdescriptor correlations, including the near-perfect anticorrelation $\rho_S(\Delta E_{\text{cd}}, R_{\text{HL}}) \approx -0.978$, which we show to be a robust geometric trend for spectra dominated by a common emission-envelope asymmetry. Benchmarked against phenomenological simulations, Hamiltonian diagonalization, and the measured MoSe₂/WSe₂ descriptor correlations, the framework turns descriptor maps into a quantitative, peak-decomposition-free probe of slow disorder and local traps in moiré excitons and, more broadly, in disordered semiconductor emitters.

DOI: [10.1103/jt25-c8fp](https://doi.org/10.1103/jt25-c8fp)

I. INTRODUCTION

Moiré heterostructures formed by vertically stacking two-dimensional (2D) materials with a small twist angle or lattice mismatch have become a central platform for engineering quantum states of electrons and excitons [1–8]. In transition-metal dichalcogenide (TMD) heterobilayers—building on the strong photoluminescence of monolayer TMDs [9–11] and their rich optical and valley physics [12–14]—the long-wavelength moiré potential reorganizes interlayer excitons into localized minibands, enabling moiré-trapped exciton physics including moiré Bose-Hubbard models, exciton crystals, and quantum emitter arrays [15–25].

Among these systems, MoSe₂/WSe₂ heterobilayers are particularly rich platforms for studying moiré-modulated optical responses. At low temperatures, the interlayer exciton photoluminescence (PL) is characteristically complex: A broad emission envelope near 1.25–1.40 eV coexists with dense, narrow spectral peaks whose microscopic assignment remains debated, with proposed contributions from moiré-trapped interlayer excitons, phonon-assisted momentum-indirect emission, donor-acceptor pair excitons, and disorder-localized states [17,26–34].

A key observation motivating the present work is that the spectral complexity is *spatially organized*. Figure 1 sketches the system: a twisted MoSe₂/WSe₂ moiré heterobilayer [Fig. 1(a)], its hyperspectral PL map [Fig. 1(b)], and the multiscale disorder landscape underlying our model [Fig. 1(c)].

Recent hyperspectral PL mapping experiments by Ahmad *et al.* [35] on a MoSe₂/WSe₂ heterobilayer reveal that nine peak-decomposition-free spectral descriptors (centroid energy E_{cent} , dominant-peak energy E_{dom} , centroid-dominant offset ΔE_{cd} , quantile width W_{80} , high/low ratio R_{HL} , sharp fraction F_S , roughness R_1 , spectral entropy S_{spec} , and integrated intensity I_{tot}) show characteristic micron-scale spatial correlations. Featurewise spatial correlation analysis yields $1/e$ decay lengths of 1.27–2.05 μm for all descriptors, exceeding the 0.85 μm optical spot size. Gaussian mixture clustering on the principal component analysis (PCA)-projected descriptor space identifies three dominant spectral families that form contiguous real-space domains. At the same time, individual pixels retain a dense multipeak structure that varies across the map, implying an unresolved local spectral manifold below optical resolution.

Here, the *local spectral manifold* denotes the optically unresolved ensemble of local emitting states—including moiré-localized and trap-related exciton states—that collectively form the PL spectrum within a single optical spot; we use the term in this physical sense, not in the differential-geometric one. More broadly, hyperspectral PL imaging is emerging as a spatially resolved probe of strain and disorder in TMD systems [36].

This combination of phenomena—resolved micron-scale spectral landscape coexisting with unresolved

*Contact author: WAKABAYASHI.Katsunori@nims.go.jp

Published by the American Physical Society under the terms of the *Creative Commons Attribution 4.0 International* license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.




FIG. 1. Overview of the system, measurement, and theoretical framework. (a) MoSe₂/WSe₂ heterobilayer with interlayer distance ~ 0.7 nm and twist angle θ , producing a long-wavelength moiré superlattice of period a_M . (b) Hyperspectral photoluminescence mapping yields a spatially resolved PL spectrum at each map position (x_0, y_0) , from which peak-decomposition-free spectral descriptors (E_{cent} , E_{dom} , ΔE_{cd} , W_{80} , R_{HL} , F_s , R_1 , S_{spec} , I_{tot}) are extracted without microscopic line assignment. (c) Theoretical multiscale disorder landscape: A slowly varying background potential $V_{\text{slow}}(\mathbf{r})$ at the micron scale combined with localized traps $V_{\text{trap}}(\mathbf{r})$ at suboptical length scales yields the micron-scale effective potential $V_{\text{eff}}(\mathbf{r}) \simeq V_{\text{slow}}(\mathbf{r}) + V_{\text{trap}}(\mathbf{r})$ that organizes the spectral descriptors hierarchically; the moiré term in the full potential [Eq. (2)] is coarse grained over the optical spot and enters only through the renormalized mass and local density of states.

local-spectral-manifold complexity—suggests that a single disorder scale is insufficient. The central experimental conclusion is that the emission is organized *hierarchically*: A slowly varying micron-scale background shapes the envelope, while local-spectral-manifold degrees of freedom produce dense fine structure within that background.

Despite extensive experimental and numerical studies of moiré exciton physics, a minimal theoretical framework for hierarchical spectral inhomogeneity as an organizing principle has not been established. Previous theoretical work has focused primarily on (1) single-particle moiré miniband structures [37–41], (2) correlated exciton and electron phases [19,20,42–44], (3) isolated defect-trapped excitons as quantum emitters [45–47], or (4) Gaussian disorder models for energy broadening without spatial structure [28]. These approaches do not directly address the *separation of scales* observed in hyperspectral mapping.

In this paper, we develop a minimal theory for the spatial organization of spectral descriptors in moiré exciton PL—minimal in the sense of containing the fewest ingredients needed to reproduce the observed descriptor hierarchy, not minimal in a microscopic sense—built around a single organizing principle: Different descriptors act as filters for different components of the disorder landscape. A descriptor that integrates spectral weight (centroid energy) averages over local trap fluctuations and tracks the smooth, micron-scale potential V_{slow} , while a descriptor that locates spectral peaks (dominant-peak energy) is additionally randomized by trap-level switching within the optical spot. This differential filtering is why the descriptor-specific correlation lengths differ, why the interdescriptor Spearman matrix has the

structure it does, and why three spectral families emerge from a continuous disorder landscape. It thereby connects the leading effective disorder parameters (ξ_s , W_s , W_t) to the observable descriptor covariance structure, without any microscopic peak assignment.

The principal contributions of this work are the following. (1) We formalize the descriptor-based disorder-filter principle introduced above into a quantitative mapping from the descriptor covariance matrix to the disorder parameters (ξ_s , W_s , W_t) [48]. (2) We derive a correlation-length hierarchy $\xi(E_{\text{cent}}) \geq \xi(E_{\text{dom}})$ from the decomposition $E_{\text{dom}} = \bar{V}_s + \delta E_{\text{dom}}$. The trap-switching fluctuation δE_{dom} adds short-range variance, which systematically shortens the correlation length of the dominant-peak energy relative to that of the centroid. (3) We show that the nontrivial content of this hierarchy lies not in the mere fact that the centroid is smoother than the peak position (expected of any average), but in its magnitude: The ratio $\xi(E_{\text{cent}})/\xi(E_{\text{dom}})$ quantitatively encodes the relative strength of the trap and slow disorder and the trap density. We turn this into an explicit, peak-decomposition-free protocol (Table V) that constrains the leading effective disorder parameters from four measured descriptor observables—the centroid correlation length, the centroid variance, the $E_{\text{cent}}-E_{\text{dom}}$ correlation, and the mean sharp fraction (the last fixing the trap density only after calibration of the sharp-fraction scale)—applicable, in principle, to hyperspectral PL maps containing both smooth envelopes and unresolved local fine structure. (4) We derive analytically the sign and approximate magnitude of the principal interdescriptor Spearman correlations. This establishes the $\Delta E_{\text{cd}}-R_{\text{HL}}$ anticorrelation as a robust spectral-shape relation,

TABLE I. Three disorder components and their characteristic length scales. Here, σ_t is an effective trap width spanning the experimentally debated sharp-emitter origins (atomic defects, strain, or electrostatic disorder), not a microscopic defect size; the theory requires only $\sigma_t \lesssim \sigma_{\text{opt}} \ll \xi_s$.

Component	Length scale	Physical origin	Role
$V_{\text{moiré}}$	$a_M \sim 10\text{--}20\text{ nm}$	Moiré superlattice	Miniband formation
V_{slow}	$\xi_s \sim 1\text{--}2\text{ }\mu\text{m}$	Strain, twist-angle gradient	Spectral domain formation
V_{trap}	$\sigma_t \lesssim \sigma_{\text{opt}}$ (few–100 nm effective)	Localized traps (defects/strain/charge)	Sharp-line generation

valid for a broad class of spectra with a dominant unimodal envelope, and shows the F_S - R_1 covariation to follow from both descriptors measuring the same trap-peak density. (5) We provide an independent model-level check via Hamiltonian diagonalization—without relying on the synthetic two-fluid spectrum construction—recovering the correlation-length hierarchy from the eigenvalue spectrum alone. (6) We construct a disorder parameter map that identifies four qualitatively distinct regimes and places the experimental system in the hierarchically disordered regime where both slow and trap disorder exceed the homogeneous linewidth.

Throughout, the theoretical framework is a *spectral-statistical organization theory*: The Hamiltonian defines a continuum spectral manifold whose coarse-grained statistics are resolved by the optical spot, and the descriptors are the observables through which the disorder hierarchy is encoded and constrained. This is distinct from a quantum-transport or localization theory—The relevant physics operates at the $\sim\mu\text{m}$ scale of the optical envelope, not at the nanometer scale of individual eigenstates.

The remainder of this paper is organized as follows. Section II introduces the effective Hamiltonian and disorder model. Section III develops the Green's function formalism and local PL spectrum. Section IV derives the statistical properties of spectral descriptors as disorder filters. Section V derives the correlation-length hierarchy relation. Section VI derives interdescriptor Spearman correlations. Section VII maps the four disorder regimes. Section VIII presents the local-density-of-states (LDOS) two-fluid structure. Section IX interprets spectral families as correlated disorder domains, and Sec. X characterizes the multiscale spectral organization that emerges in the hierarchical regime. Section XI describes Hamiltonian diagonalization and phenomenological simulations. Section XII discusses limitations and parameter extraction. Section XIII concludes.

II. MODEL: EFFECTIVE EXCITON HAMILTONIAN

This section constructs the effective Hamiltonian on which the rest of the paper rests. We model the interlayer exciton as a single center-of-mass (COM) particle moving in an effective potential, and assemble that potential from three physically distinct ingredients—the moiré superlattice, a slowly varying long-wavelength disorder, and sparse local traps—each introduced in turn below. We then identify the separation of length scales among these ingredients, which is what makes the descriptor-filter analysis of the following sections tractable.

A. Center-of-mass Hamiltonian

We describe the interlayer exciton by its COM coordinate $\mathbf{r}$ in the moiré plane. The internal exciton structure is integrated out, yielding an effective single-particle problem for the COM motion:

$$H = -\frac{\hbar^2}{2M}\nabla^2 + V_{\text{eff}}(\mathbf{r}), \quad (1)$$

where M is the interlayer exciton effective mass. For $\text{MoSe}_2/\text{WSe}_2$, the electron and hole are spatially separated into different layers [21,23], giving $M \approx m_e + m_h \approx 1.0\text{--}1.5m_0$ [12,49,50], where m_0 is the free-electron mass.

The effective potential is decomposed into three components acting on different length scales:

$$V_{\text{eff}}(\mathbf{r}) = V_{\text{moiré}}(\mathbf{r}) + V_{\text{slow}}(\mathbf{r}) + V_{\text{trap}}(\mathbf{r}). \quad (2)$$

These three components describe physics at three distinct length scales, as summarized in Table I.

B. Moiré potential

The moiré superlattice potential is modeled by the lowest-order hexagonally symmetric Fourier components:

$$V_{\text{moiré}}(\mathbf{r}) = 2V_m \sum_{j=1}^3 \cos(\mathbf{G}_j \cdot \mathbf{r} + \phi_j), \quad (3)$$

where $\mathbf{G}_j$ are the three inequivalent moiré reciprocal lattice vectors satisfying $|\mathbf{G}_j| = 4\pi/(a_M\sqrt{3})$, V_m controls the modulation amplitude, and ϕ_j encodes the local stacking registry. For small twist angle θ and lattice constant a_0 , the moiré period is $a_M \approx a_0/\theta$.

For a twist angle $\theta \sim 2^\circ\text{--}4^\circ$, $a_M \sim 10\text{--}20\text{ nm}$, and the moiré potential depth is estimated at $V_m \sim 10\text{--}50\text{ meV}$ from first-principles calculations [37–40,51].

This term confines excitons to moiré unit cells, creating localized minibandlike states on the nanometer scale. For the purposes of spatial PL statistics at the micron scale, the moiré potential can be treated as a fast-varying background that renormalizes the exciton effective mass and density of states (DOS), but does not directly set the micron-scale spatial correlations.

C. Long-wavelength correlated disorder

The dominant component governing micron-scale spectral organization is a slowly varying, spatially correlated random potential:

$$\langle V_{\text{slow}}(\mathbf{r}) \rangle = 0, \quad \langle V_{\text{slow}}(\mathbf{r}) V_{\text{slow}}(\mathbf{r}') \rangle = W_s^2 \mathcal{C}_s(|\mathbf{r} - \mathbf{r}'|), \quad (4)$$

with Gaussian correlation kernel:

$$C_s(r) = \exp\left(-\frac{r^2}{2\xi_s^2}\right). \quad (5)$$

The characteristic parameters are $W_s \sim 5\text{--}20$ meV (disorder amplitude) and $\xi_s \sim 1\text{--}2$ μm (correlation length), much larger than the moiré period a_M .

The physical origins of V_{slow} are diverse. Spatial variations in the local twist angle $\delta\theta(\mathbf{r})$, directly imaged in twisted TMD bilayers [52,53], shift the interlayer exciton energy as $\partial E/\partial\theta \times \delta\theta$ (the twist-angle dependence of interlayer-exciton properties in TMD heterobilayers is well established [29]). Inhomogeneous strain fields $\varepsilon_{ij}(\mathbf{r})$ couple through the deformation potential $\sim D_{ij}\varepsilon_{ij}$ [54]. Electrostatic potential fluctuations arise from trapped charges or dielectric inhomogeneity [55]. Finally, moiré reconstruction domain walls introduce steplike modulations [52,53,56]. All of these generically produce smooth fields with correlation lengths set by the macroscopic sample quality, typically 1–5 μm in high-quality van der Waals assemblies.

We work in the Gaussian approximation (4) as the minimal model, realized numerically by the spectral [fast Fourier transform (FFT)] method of Appendix A; extensions to non-Gaussian slow disorder are discussed in Sec. XII.

D. Local trap disorder

Local trap potentials are introduced as a sum of attractive Gaussians centered at random positions $\{\mathbf{R}_i\}$:

$$V_{\text{trap}}(\mathbf{r}) = \sum_i u_i \exp\left(-\frac{|\mathbf{r} - \mathbf{R}_i|^2}{2\sigma_t^2}\right), \quad (6)$$

where $\mathbf{R}_i$ are drawn from a Poisson point process with density n_t ; $u_i \leq 0$ are random trap depths drawn from a distribution $P(u)$ with mean $\bar{u} < 0$ and characteristic depth scale W_t [the simulations realize $P(u)$ as a half-normal, $u_i = -|\mathcal{N}(0, W_t)|$], so that W_t is the scale of the underlying normal, with second moment $\langle u^2 \rangle = W_t^2$ and variance $W_t^2(1 - 2/\pi)$; and $\sigma_t \ll \xi_s$ is the trap radius.

Each trap acts as a local potential minimum that can bind an exciton and produce a sharp spectral line. The trap correlation function is

$$\langle V_{\text{trap}}(\mathbf{r}) V_{\text{trap}}(\mathbf{r}') \rangle = n_t \langle u^2 \rangle \times \pi \sigma_t^2 \exp\left(-\frac{|\mathbf{r} - \mathbf{r}'|^2}{4\sigma_t^2}\right), \quad (7)$$

with correlation length $\sim \sqrt{2}\sigma_t \ll \xi_s$. This spatial white-noise-like character of the trap disorder is what distinguishes it from V_{slow} . Microscopic sources of such localized potentials include native point defects [45,57], strained nanobubbles and strain-localized emitters [31,47], and adsorbate-induced potential wells.

E. Parameter hierarchy and scale separation

The effective model assumes a clear separation of length scales,

$$\sigma_t, a_M \ll \xi_s \ll L, \quad (8)$$

where L is the sample (or map) size. In the experiment [35], $L \sim 8$ μm and $\xi_s \sim 2$ μm . The two microscopic scales of the model lie far below ξ_s : the moiré period $a_M \sim 10\text{--}20$ nm and

the trap radius σ_t , which ranges from the atomic/defect scale (a few nanometer) up to ~ 100 nm in our robustness tests.

We deliberately leave the relative ordering of σ_t and a_M unspecified. Instead, σ_t is treated as an *effective* potential width rather than a microscopic trap dimension, because the origin of the sharp emitters—moiré-trapped excitons, atomic defects, or strain and electrostatic disorder—remains experimentally debated [26,28]. These possible origins have characteristic sizes ranging from a few nm up to ~ 100 nm, so no single microscopic value of σ_t is preferred.

The theory requires only $\sigma_t \lesssim \sigma_{\text{opt}} \ll \xi_s$ (with σ_{opt} the optical-spot radius), so that each trap acts as a subspot localized emitter, with many such emitters contributing to a single measured spectrum. With this separation, σ_t enters the descriptor statistics only through the product $n_t W_t^2 \sigma_t^2$. This combination is not arbitrary: It is the variance of the trap potential V_{trap} . Summing the squared contribution of each well over the Poisson-distributed trap positions gives the shot-noise variance (Campbell's theorem [58,59]) $\text{Var} V_{\text{trap}} \propto n_t W_t^2 \sigma_t^2$, with the three factors being the areal density n_t , the depth variance W_t^2 , and the well area $\sim \sigma_t^2$. We identify it as the effective trap disorder Γ_t in Sec. VII, and the descriptor variances that depend on it are derived in Appendix C.

Because σ_t appears only inside this product, its value is not pinned down on its own: Any change in σ_t can be offset by a compensating change in the trap density ($n_t \propto \sigma_t^{-2}$), or in the depth variance W_t^2 , that keeps the product—and hence the leading trap-loading contribution to the descriptor variances—approximately unchanged. In other words, σ_t , n_t , and W_t^2 are *degenerate*: They enter only in this one combination, so the predictions are insensitive to the precise value of σ_t within the allowed window. The only place σ_t sets an independent length scale is the hierarchy floor $\xi_t \equiv \sqrt{2}\sigma_t$ of Eq. (38); even for $\sigma_t \sim 100$ nm, this remains well below ξ_s ($\xi_t/\xi_s \sim 0.07$) and therefore does not bind the hierarchy. The correlation-length hierarchy and the interdescriptor correlations are thus controlled primarily by (ξ_s, W_s, W_t, n_t) , with only weak dependence on σ_t within the suboptical range; a diagonalization sweep over $\sigma_t = 50\text{--}125$ nm confirms this explicitly (Appendix D).

This separation (8) justifies treating the disorder components as statistically independent and analytically tractable at each scale. Figure 2 shows a representative realization of the resulting hierarchical disorder landscape: the slow potential V_{slow} [Fig. 2(a)], the local trap potential V_{trap} [Fig. 2(b)], and the effective potential $V_{\text{eff}} = V_{\text{slow}} + V_{\text{trap}}$ [Fig. 2(c)].

III. GREEN'S FUNCTION FORMALISM AND LOCAL PL SPECTRUM

The spectral descriptors introduced above are all computed from the *local* PL spectrum $I(E, \mathbf{R})$ —the emission collected from an optical spot centered at position $\mathbf{R}$. The goal of this section is to express that observable directly in terms of the disorder Hamiltonian H of Sec. II, so that the statistics of the descriptors can later be traced back to the underlying disorder parameters.

The retarded Green's function is the natural tool for expressing $I(E, \mathbf{R})$ in terms of H . It resolves H into its spectral content, and its imaginary part yields the LDOS: the




FIG. 2. Hierarchical disorder landscape. (a) Slow correlated disorder $V_{\text{slow}}(\mathbf{r})$ with amplitude $W_s = 6$ meV and correlation length $\xi_s = 2$ μm , generated by FFT-based Gaussian filtering. The smooth red/blue domains of size $\sim \xi_s$ represent the slow-disorder component that gives rise to micron-scale spectral families in the model. (b) Trap disorder $V_{\text{trap}}(\mathbf{r})$ consisting of $N_t = 133$ Gaussian potential wells ($W_t = 12$ meV, $\sigma_t = 50$ nm, $n_t = 2$ μm^{-2} , consistent with $n_t L^2 \approx 128$ over the 8×8 μm^2 domain). The color scale is saturated at the second percentile of V_{trap} for visibility; because the field is near zero between the sparse wells, this displayed lower bound is much shallower than the deepest individual well centers, whose bare trap-depth distribution is parametrized by $W_t = 12$ meV. Red crosses mark the trap centers, which are drawn from $P(\mathbf{r}) \propto \exp(-V_{\text{slow}}(\mathbf{r})/E_{\text{bias}})$, with $E_{\text{bias}} = 7$ meV, biasing the trap positions, or optically active trap sites, toward slow-potential minima. (c) Effective potential $V_{\text{eff}} = V_{\text{slow}} + V_{\text{trap}}$, showing the combined landscape experienced by excitons. The gold dashed rectangle indicates the 20×20 measurement grid area (spanning 7.6×7.6 μm^2 , i.e., a 0.4 μm pitch matching the experimental 400 nm step, with a 0.2 μm margin from each edge).

spectral weight that the eigenstates of H place at energy E and position $\mathbf{r}$. The measured local spectrum is then this LDOS filtered by the finite optical spot and weighted by the radiative coupling of each eigenstate to light. In this way, the formalism converts the eigenstates $\{\psi_n, E_n\}$, which encode the multiscale disorder, into the observable $I(E, \mathbf{R})$ without fitting or identifying individual spectral peaks—the peak-decomposition-free approach on which the descriptor analysis is based. We build this chain in three steps: the Green’s function, the LDOS it generates, and the radiatively weighted, optically filtered local spectrum.

A. Retarded Green’s function

The retarded Green’s function is the resolvent of the exciton Hamiltonian: Its poles sit at the eigenenergies E_n and its residues carry the eigenfunctions ψ_n , so it packages the entire spectral content of H in a single object [60],

$$G(\mathbf{r}, \mathbf{r}'; E) = \langle \mathbf{r} | (E + i\eta - H)^{-1} | \mathbf{r}' \rangle = \sum_n \frac{\psi_n(\mathbf{r}) \psi_n^*(\mathbf{r}')}{E - E_n + i\eta}, \quad (9)$$

where $\psi_n(\mathbf{r})$ and E_n are the eigenstates and eigenenergies of H , and $\eta \rightarrow 0^+$ is an infinitesimal broadening.

B. Local density of states

The LDOS is

$$\rho(E, \mathbf{r}) = -\frac{1}{\pi} \text{Im} G(\mathbf{r}, \mathbf{r}; E) = \sum_n |\psi_n(\mathbf{r})|^2 \delta(E - E_n). \quad (10)$$

In practice, the delta function is replaced by a Lorentzian or Gaussian with linewidth $\eta \sim 1\text{--}3$ meV (phenomenological

exciton lifetime broadening):

$$\rho_\eta(E, \mathbf{r}) = \sum_n |\psi_n(\mathbf{r})|^2 g_\eta(E - E_n), \quad g_\eta(E) = \frac{\eta/\pi}{E^2 + \eta^2}. \quad (11)$$

C. Local PL spectrum

The connection to experiment is set by the radiative coupling of the COM exciton eigenstates to light. Because the internal $1s$ relative wave function is essentially rigid, the light couples to each COM eigenstate $\psi_n(\mathbf{r})$ through its *coherent* dipole (oscillator-strength) matrix element, the amplitude integral of the COM wave function over the emission region, so that the far-field radiative weight of eigenstate n is $\propto |\int \psi_n d^2\mathbf{r}|^2$ rather than the incoherent probability density $|\psi_n|^2$; this is the standard result for disorder-localized exciton emission in quantum wells and related nanostructures [61–63]. For a local measurement that collects from an optical spot centered at $\mathbf{R}$ with amplitude collection mode W_{amp} , the locally detected PL spectrum is

$$I(E, \mathbf{R}) = \sum_n f_{\text{occ}}(E_n) A_n(\mathbf{R}) g_\eta(E - E_n), \quad (12)$$

with the coherent radiative weight

$$A_n(\mathbf{R}) = |M_n(\mathbf{R})|^2, \quad M_n(\mathbf{R}) = \int W_{\text{amp}}(\mathbf{r} - \mathbf{R}) \psi_n(\mathbf{r}) d^2\mathbf{r}, \quad (13)$$

where W_{amp} is the (real, positive) collection-mode amplitude and $W_{\text{opt}} = |W_{\text{amp}}|^2$ is the intensity point-spread function. For a diffraction-limited Gaussian spot, both are

Gaussian,

$$W_{\text{amp}}(\mathbf{r}) = \frac{1}{\sqrt{2\pi\sigma_{\text{opt}}^2}} \exp\left(-\frac{|\mathbf{r}|^2}{4\sigma_{\text{opt}}^2}\right), \quad (14)$$

$$W_{\text{opt}}(\mathbf{r}) = |W_{\text{amp}}(\mathbf{r})|^2 = \frac{1}{2\pi\sigma_{\text{opt}}^2} \exp\left(-\frac{|\mathbf{r}|^2}{2\sigma_{\text{opt}}^2}\right), \quad (15)$$

so that W_{amp} has width $\sqrt{2}\sigma_{\text{opt}}$ and the normalized intensity PSF W_{opt} has width $\sigma_{\text{opt}} = \text{FWHM}/(2\sqrt{2\ln 2}) \simeq 0.36 \mu\text{m}$ [experimental full width at half maximum (FWHM) = $0.85 \mu\text{m}$]. The occupation factor is $f_{\text{occ}}(E)$. Only the coarse-graining scale σ_{opt} enters the descriptor correlations; the overall normalization sets only the absolute emission intensity. The coherent weight (13) preferentially enhances nodeless localized states and suppresses spatially extended states with sign-changing wave functions. The mechanism is the cancellation in the amplitude integral $M_n = \int W_{\text{amp}} \psi_n$: A nodeless, well-localized trap state keeps a single sign across its small extent, so its contributions add up coherently and A_n is large; an extended state whose wave function oscillates in sign has its positive and negative contributions largely cancel, so A_n is small. Thus, in this approximation, localized traplike states carry larger radiative weights, whereas extended background states tend to be suppressed by phase cancellation—a PL-level analog of the mobility-edge selectivity known for disordered quantum-well excitons [61,62].

For the analytic descriptor theory, it is convenient to use the incoherent *local-density approximation*, in which $|M_n(\mathbf{R})|^2$ is replaced by $\int d^2\mathbf{r} W_{\text{opt}}(\mathbf{r} - \mathbf{R}) |\psi_n(\mathbf{r})|^2$, so that Eq. (12) becomes the optically filtered LDOS:

$$I(E, \mathbf{R}) = \int d^2\mathbf{r} W_{\text{opt}}(\mathbf{r} - \mathbf{R}) f_{\text{occ}}(E) \rho_\eta(E, \mathbf{r}). \quad (16)$$

Because both weights are localized within the optical spot, the descriptor covariance and the correlation-length hierarchy derived below are essentially the same for the coherent and incoherent forms. We therefore use the transparent LDOS form (16) both in the analytic treatment and in the microscopic diagonalization of Sec. XID and verify in Appendix E that replacing it by the microscopically motivated coherent radiative form (13) leaves the descriptor statistics unchanged within the numerical seed-to-seed uncertainty (the two weightings differ by ≤ 0.01 for every descriptor).

Approximating $f_{\text{occ}} \approx 1$ for all states in the analyzed PL window yields the working form

$$I(E, \mathbf{R}) = \sum_n A_n(\mathbf{R}) g_\eta(E - E_n). \quad (17)$$

Here, $A_n(\mathbf{R})$ denotes the local optical weight in either representation: the coherent radiative form of Eq. (13) or the incoherent local-density approximation $A_n(\mathbf{R}) = \int d^2\mathbf{r} W_{\text{opt}}(\mathbf{r} - \mathbf{R}) |\psi_n(\mathbf{r})|^2$ [Eq. (16)] used in the analytic treatment below. A full treatment would fix f_{occ} from the relaxation kinetics between disorder eigenstates [62], which at low temperature can be strongly energy selective—for example, through a mobility edge—and which we do not attempt to model here. Instead, we adopt the constant-occupation form $f_{\text{occ}} \simeq 1$ as a deliberate minimal (null) model, *not* as a thermal-equilibrium approximation, so that the descriptor-correlation hierarchy is

attributed to the disorder-filtering mechanism rather than to occupation-induced weighting. Within this minimal model, a spatially uniform, smooth energy-dependent occupation would mainly reshape the common spectral envelope over the analyzed window ($\sim 150 \text{ meV}$, i.e., $1.25\text{--}1.40 \text{ eV}$) and is not expected to change the spatial correlation hierarchy, since the latter is set by the spatial statistics of the eigenenergies E_n and weights $A_n(\mathbf{R})$. A strongly energy-selective occupation, or mobility-edge filtering of the kind emphasized in Refs. [61,62], would constitute an additional physical weighting beyond the present minimal model, which we note here as a limitation. The optical spot intensity function is modeled as

$$W_{\text{opt}}(\mathbf{r}) = \frac{1}{2\pi\sigma_{\text{opt}}^2} \exp\left(-\frac{|\mathbf{r}|^2}{2\sigma_{\text{opt}}^2}\right), \quad (18)$$

with $\sigma_{\text{opt}} = \text{FWHM}/(2\sqrt{2\ln 2}) \approx 0.36 \mu\text{m}$ for an FWHM of the optical spot of $0.85 \mu\text{m}$.

D. Spectral decomposition: Background and trap contributions

A key analytical insight is to decompose $I(E, \mathbf{R})$ into contributions from states with different spatial extents. We classify eigenstates using the inverse participation ratio (IPR), which provides a standard measure of real-space localization [61] and thus naturally separates spatially extended background states (small IPR) from trap-localized states (large IPR). We define the localization length $\ell_n = \text{IPR}_n^{-1/2}$, where $\text{IPR}_n = \int d^2\mathbf{r} |\psi_n(\mathbf{r})|^4$ (with ψ_n normalized). The IPR works as an inverse effective area: For a state spread uniformly over an area A , one has $|\psi_n|^2 \sim 1/A$ and $\text{IPR}_n \sim 1/A$, so $\ell_n = \text{IPR}_n^{-1/2}$ measures the linear extent over which the state has appreciable weight—small ℓ_n for tightly localized trap states and large ℓ_n for extended background states [61]. A cutoff ℓ^* (of order the trap size) then separates the two families:

$$I(E, \mathbf{R}) = I_{\text{bg}}(E, \mathbf{R}) + I_{\text{sharp}}(E, \mathbf{R}), \quad (19)$$

$$I_{\text{bg}}(E, \mathbf{R}) = \sum_{n: \ell_n > \ell^*} A_n(\mathbf{R}) g_\eta(E - E_n), \quad (20)$$

$$I_{\text{sharp}}(E, \mathbf{R}) = \sum_{n: \ell_n \leq \ell^*} A_n(\mathbf{R}) g_\eta(E - E_n), \quad (21)$$

where ℓ^* is a crossover localization length (see Sec. VIII). States with $\ell_n > \ell^*$ are spatially extended on the scale of the optical spot and contribute to the smooth background envelope I_{bg} . Trap-concentrated states ($\ell_n \leq \ell^*$) contribute narrow lines to I_{sharp} . The IPR here serves to define this conceptual two-component decomposition, with ℓ^* acting as a practical crossover rather than a sharp microscopic phase boundary. In the strongly disordered regime of the numerical simulations, where the IPR is set by the disorder landscape rather than by quantum interference, we apply an equivalent energy-based classification of the same two families (Sec. VIII).

E. Optical coarse graining of slow disorder

The slow-disorder component enters the measured PL spectrum primarily through optical coarse graining. The experimentally measured spectrum at position $\mathbf{R}$ is not the LDOS at a single microscopic point, but the LDOS averaged

over the optical point-spread function $W_{\text{opt}}(\mathbf{r} - \mathbf{R})$ [Eq. (17)]. Since $V_{\text{slow}}(\mathbf{r})$ varies on the length scale ξ_s , much larger than the optical spot size σ_{opt} ($\sigma_{\text{opt}}/\xi_s \ll 1$), the smooth part of the local spectral manifold within one optical spot is shifted almost uniformly by the local average of V_{slow} .

We therefore define the optically averaged slow potential as

$$\bar{V}_s(\mathbf{R}) = \int d^2\mathbf{r} W_{\text{opt}}(\mathbf{r} - \mathbf{R}) V_{\text{slow}}(\mathbf{r}). \quad (22)$$

This quantity is the energy shift of the background PL envelope observed at the optical position $\mathbf{R}$, rather than the first-order energy correction of a single microscopic eigenstate; it requires no assumption about the detailed profile of individual wave functions.

Equivalently, for the background contribution to the LDOS, one may write

$$\rho_{\text{bg}}(E, \mathbf{R}) \simeq \rho_{\text{bg}}^{(0)}(E - \bar{V}_s(\mathbf{R})), \quad (23)$$

up to corrections controlled by the small ratio $\sigma_{\text{opt}}/\xi_s$. This relation expresses the central coarse-graining assumption used below: The smooth background within one optical spot is rigidly displaced by $\bar{V}_s(\mathbf{R})$, so that centroidlike descriptors, which integrate spectral weight over the smooth envelope, primarily track $\bar{V}_s(\mathbf{R})$, whereas peak-position descriptors remain sensitive to additional local trap-level fluctuations within the same optical spot.

IV. STATISTICAL THEORY OF SPECTRAL DESCRIPTORS

This section makes the disorder-filter principle quantitative. We first define the spectral descriptors—simple, peak-decomposition-free functionals of the local spectrum—and then derive the statistics of the most important ones, showing explicitly that the centroid energy tracks the slow disorder, while the spectral roughness and entropy track the traps. These results are the analytical basis for the correlation-length hierarchy and the interdescriptor correlations derived in the following sections.

A. Definitions of spectral descriptors

We define the full set of spectral descriptors following the operational conventions of the experiment [35], so that simulated and measured descriptors are directly comparable. Each descriptor is a simple functional of the raw local spectrum $I(E, \mathbf{R})$ that requires no peak fitting or microscopic line assignment. The set is deliberately chosen so that different descriptors are sensitive to different aspects of the spectral shape—its overall position, its width, its asymmetry about the dominant peak, and its fine structure—and therefore, through the disorder-filter principle, to different components of the underlying disorder: Descriptors that integrate the whole spectrum (such as the centroid) track the smooth background, whereas peak- and roughness-based descriptors respond to the sharp, trap-generated fine structure. We introduce them in turn, in each case stating what the descriptor measures before giving its definition.

1. Centroid energy

The centroid is the intensity-weighted mean emission energy—the first moment of the local spectrum. Because it integrates over the entire spectrum, it averages out the subspot trap fluctuations and tracks the smooth, slowly varying background set by V_{slow} ; it is therefore the descriptor most directly tied to the micron-scale disorder (Sec. V). It is defined by

$$E_{\text{cent}}(\mathbf{R}) = \frac{\int E I(E, \mathbf{R}) dE}{\int I(E, \mathbf{R}) dE}. \quad (24)$$

2. Dominant-peak energy

The dominant-peak energy is the position of the tallest spectral feature—the mode of the local spectrum. In contrast to the centroid, it is set by whichever single state carries the most optical weight within the optical spot, so it responds to discrete trap-level switching as the spot is moved and, as shown below, acquires a shorter correlation length than E_{cent} . It is defined by

$$E_{\text{dom}}(\mathbf{R}) = \arg \max_E I(E, \mathbf{R}). \quad (25)$$

Here, $\arg \max_E$ denotes the energy at which the local PL spectrum reaches its maximum; for discretized spectra, it is taken as the energy bin with the largest intensity.

3. Centroid-dominant offset

The offset between the centroid and the dominant peak measures how far the intensity-weighted mean lies from the mode, i.e., the *skewness* of the emission envelope about its own maximum. It vanishes for a symmetric line and takes a definite sign as spectral weight is redistributed to the low- or high-energy side of the dominant peak, making it a sensitive, peak-decomposition-free probe of envelope asymmetry. It is defined by

$$\Delta E_{\text{cd}}(\mathbf{R}) = E_{\text{cent}}(\mathbf{R}) - E_{\text{dom}}(\mathbf{R}). \quad (26)$$

4. Quantile width

The quantile width characterizes the overall energy spread of the emission without reference to any individual peak. Using the central-80% band rather than a variance or a fitted linewidth makes it robust to weak far tails and to the multipeak structure of the local spectrum, so that it reflects the width of the dominant emission band. Specifically, $W_{80}(\mathbf{R})$ is the energy width of the band containing the central 80% of the integrated PL intensity, i.e., the interval between the 10th and 90th cumulative-intensity percentiles.

5. High/low ratio

The high/low ratio quantifies the asymmetry of the emitted intensity about the dominant peak—how much spectral weight lies below the main line relative to above it. It captures the same envelope skew as the offset ΔE_{cd} , but through the distribution of intensity rather than of energy, which is why the two turn out to be so tightly linked (Sec. VI). Following the experimental convention [35], the spectrum is split at the

dominant peak E_{dom} (not at a fixed global energy):

$$R_{\text{HL}}(\mathbf{R}) = \frac{\int_{-\infty}^{E_{\text{dom}}} I(E, \mathbf{R}) dE}{\int_{E_{\text{dom}}}^{\infty} I(E, \mathbf{R}) dE}, \quad (27)$$

so that $R_{\text{HL}} > 1$ when spectral weight accumulates below the dominant emission peak. This shared E_{dom} reference is essential: It is what makes R_{HL} and ΔE_{cd} two measures of the same envelope asymmetry about the dominant peak, and hence tightly anticorrelated. Referencing the split to E_{cent} or to a fixed spectral quantile instead measures a different quantity and weakens or reverses the relation, as expected.

6. Sharp fraction

The sharp fraction is the fraction of the emitted intensity carried by narrow, trap-related lines rather than by the broad background envelope. It is a direct, peak-decomposition-free proxy for the local trap loading—the number and depth of sharp emitters within the optical spot—and is therefore the descriptor most sensitive to the trap component of the disorder. It is defined by

$$F_S(\mathbf{R}) = \frac{I_{\text{sharp,tot}}(\mathbf{R})}{I_{\text{tot}}(\mathbf{R})}, \quad (28)$$

where $I_{\text{sharp,tot}} = \int I_{\text{sharp}}(E, \mathbf{R}) dE$ is the intensity carried by *sharp* spectral features. Operationally, sharp peaks are identified by local-maximum detection: A local maximum qualifies as sharp if its prominence exceeds a fraction p of the pixel maximum ($p = 0.02$) and its width at half-prominence is below a threshold w ($w = 10$ meV), which selects narrow trap lines while excluding the broad background envelope (FWHM ~ 35 meV); I_{sharp} is the intensity within $\pm w/2$ of the qualifying peaks. This peak-detection criterion—rather than a fixed linewidth threshold or a parametric background subtraction—is robust to its parameters: Varying p over 0.01–0.05 and w over 8–15 meV changes $\rho_S(F_S, R_1)$ by less than 0.03 (from +0.90 to +0.92), while the mean $\langle F_S \rangle$ shifts only within 0.25–0.39.

7. Spectral roughness

The spectral roughness is the total variation of the normalized spectrum—its integrated absolute slope—and so quantifies the amount of sharp structure it contains: A smooth envelope gives a small R_1 , while each additional narrow trap line adds to the total variation. It is thus a peak-detection-free proxy for the density of fine structure, complementary to the sharp fraction F_S (with which it is strongly correlated; Sec. VI). Normalizing by the integrated intensity makes it independent of the absolute emission strength:

$$R_1(\mathbf{R}) = \frac{\int |dI(E, \mathbf{R})/dE| dE}{\int I(E, \mathbf{R}) dE}. \quad (29)$$

8. Spectral entropy

The spectral entropy measures how uniformly the emitted intensity is distributed across energy, treating the normalized spectrum as a probability distribution. It is small when the emission is concentrated in a few features and approaches unity for a broad, feature-rich spectrum, and it therefore provides a peak-decomposition-free measure of local spectral

complexity that complements the quantile width. Normalized by $\ln N$ to lie in $[0, 1]$, it reads

$$S_{\text{spec}}(\mathbf{R}) = -\frac{1}{\ln N} \sum_k p_k(\mathbf{R}) \ln p_k(\mathbf{R}), \quad (30)$$

where $p_k = I(E_k, \mathbf{R}) / \sum_j I(E_j, \mathbf{R})$ is the normalized spectral weight on pixel k of the spectrometer and N is the total number of energy pixels.

The fine-structure descriptors that involve a derivative or an energy-pixel sum (R_1, S_{spec}) are not purely intrinsic quantities: Their *absolute* values depend on the spectral resolution, the energy binning, any smoothing applied before analysis, and the experimental noise floor. R_1 involves an energy derivative and so rescales under smoothing or spectral broadening; S_{spec} depends on the effective number of spectral bins and on the noise floor. We therefore do not treat these absolute values as instrument independent, and they should not be compared across datasets unless the resolution and preprocessing are matched. In our simulations, they are evaluated at the experimental spectral resolution.

What we do compare is one level removed from the absolute values: the spatial covariance of the descriptors, the interdescriptor correlation signs, and the correlation lengths, all computed with the *same* preprocessing applied to the experimental and simulated spectra. These quantities are expected to be more robust to a uniform rebinning or broadening, because the same operation acts at every spatial pixel, while the pixel-to-pixel variation is set by the underlying disorder landscape. We verified this explicitly only for F_S , whose F_S – R_1 correlation shifts by less than 0.03 under the threshold variations above; we do not claim exact invariance for R_1 or S_{spec} , and a full propagation of instrumental noise and resolution through the covariance matrix is left to future work. Accordingly, R_1 , F_S , and S_{spec} are best regarded as *comparative* descriptors defined together with a specified spectral-preprocessing protocol, rather than as absolute instrument-independent observables; consistent preprocessing of experiment and simulation is what makes the comparison meaningful.

B. Centroid energy: Slow disorder dominance

We now derive the statistical properties of each descriptor from the disorder model.

Proposition 1 (Centroid energy tracks slow disorder). When the integrated centroid is dominated by the smooth background envelope rather than by the integrated weight of the sharp trap lines,

$$E_{\text{cent}}(\mathbf{R}) \approx E_{\text{cent}}^{(0)} + \bar{V}_s(\mathbf{R}), \quad (31)$$

where $E_{\text{cent}}^{(0)}$ is the disorder-free centroid energy and $\bar{V}_s(\mathbf{R})$ is the optical-spot-averaged slow disorder (22).

Proof. From Eqs. (17) and (24),

$$E_{\text{cent}}(\mathbf{R}) = \frac{\sum_n A_n(\mathbf{R}) E_n}{\sum_n A_n(\mathbf{R})}. \quad (32)$$

Using the rigid background displacement $E_n \approx E_n^{(0)} + \bar{V}_s(\mathbf{R}_n)$ from Eq. (23), and noting that $A_n(\mathbf{R})$ is concentrated near states ψ_n whose support overlaps the optical spot at $\mathbf{R}$, so

$\bar{V}_s(\mathbf{R}_n) \approx \bar{V}_s(\mathbf{R})$ for all significantly contributing states:

$$E_{\text{cent}}(\mathbf{R}) \approx \frac{\sum_n A_n(\mathbf{R})(E_n^{(0)} + \bar{V}_s(\mathbf{R}))}{\sum_n A_n(\mathbf{R})} = E_{\text{cent}}^{(0)} + \bar{V}_s(\mathbf{R}). \quad (33)$$

Corollary 1 (Centroid correlation function).

$$C_{E_{\text{cent}}}(r) \equiv \frac{\langle \delta E_{\text{cent}}(\mathbf{R}) \delta E_{\text{cent}}(\mathbf{R}') \rangle}{\langle \delta E_{\text{cent}}^2 \rangle} \approx \exp\left(-\frac{r^2}{2(\xi_s^2 + 2\sigma_{\text{opt}}^2)}\right), \quad (34)$$

for $r = |\mathbf{R} - \mathbf{R}'|$, where $\delta E_{\text{cent}} \equiv E_{\text{cent}} - \langle E_{\text{cent}} \rangle$.

Proof. $\bar{V}_s(\mathbf{R}) = \int d^2\mathbf{r} W_{\text{opt}}(\mathbf{r} - \mathbf{R}) V_{\text{slow}}(\mathbf{r})$ is a Gaussian-filtered version of V_{slow} . Setting $r = |\mathbf{R} - \mathbf{R}'|$, the convolution of two Gaussians gives

$$\begin{aligned} \langle \bar{V}_s(\mathbf{R}) \bar{V}_s(\mathbf{R}') \rangle &= \int d^2\mathbf{r} d^2\mathbf{r}' W_{\text{opt}}(\mathbf{r} - \mathbf{R}) W_{\text{opt}}(\mathbf{r}' - \mathbf{R}') \\ &\quad \times W_s^2 e^{-|\mathbf{r} - \mathbf{r}'|^2 / 2\xi_s^2} \\ &\simeq W_s^2 \exp\left(-\frac{r^2}{2(\xi_s^2 + 2\sigma_{\text{opt}}^2)}\right), \end{aligned} \quad (35)$$

since the convolution of Gaussians with widths σ_{opt} , σ_{opt} , ξ_s yields a Gaussian with width $\sqrt{\xi_s^2 + 2\sigma_{\text{opt}}^2}$. The optical filtering reduces the prefactor to $\text{Var}(\bar{V}_s) = W_s^2 \xi_s^2 / (\xi_s^2 + 2\sigma_{\text{opt}}^2)$; for $\xi_s \gg \sigma_{\text{opt}}$, this is $\simeq W_s^2$, as used above (equivalently, one may read W_s here as the filtered amplitude $W_{s,\text{eff}} \simeq W_s$). Since the normalized correlation (34) divides by the variance, this prefactor cancels and Eq. (34) is exact regardless of the reduction. The effective correlation length is

$$\xi_{\text{eff}} = \sqrt{\xi_s^2 + 2\sigma_{\text{opt}}^2}. \quad (36)$$

Since $\xi_s \gg \sigma_{\text{opt}}$ (2 μm versus 0.36 μm), we have $\xi_{\text{eff}} \approx \xi_s$. ■

Remark 1. The optical spot smoothing is negligible when $\xi_s \gg \sigma_{\text{opt}}$, which is satisfied experimentally. The measured 1/e correlation length of E_{cent} directly yields ξ_s .

C. Spectral entropy and roughness: Trap disorder dominance

In contrast to E_{cent} , the spectral entropy and roughness primarily reflect the local spectral manifold complexity generated by trap disorder.

Proposition 2 (Entropy and roughness track trap density). For fixed slow disorder, the spectral entropy S_{spec} and roughness R_1 are increasing functions of the effective trap density $n_t^{\text{eff}}(\mathbf{R}) = n_t \cdot f(\bar{V}_s(\mathbf{R}))$, where f is an increasing function of the slow disorder background (deep traps attract excitons more effectively in a background minimum).

Justification. Each additional trap line resolved within the optical spot adds total variation to $I(E, \mathbf{R})$, raising R_1 , and adds a spectral component, raising S_{spec} toward its bound; both therefore increase with the number of trap lines per spot, i.e., with n_t^{eff} . The enhancement $f(\bar{V}_s)$ grows toward slow-potential minima, where deeper local wells bind excitons more effectively. The argument is monotonic rather than exact

(near-degenerate lines contribute subadditively), consistent with the $\propto n_t W_t^2 \sigma_t^2$ variance scaling of Appendix C.

This proposition has an important corollary: Even the trap-sensitive descriptors S_{spec} and R_1 should show micron-scale spatial correlations, because the effective trap coupling is modulated by $\bar{V}_s(\mathbf{R})$. Specifically, regions with lower E_{cent} (deeper V_{slow} minima) will show enhanced F_S and R_1 because traps in such regions create deeper bound states relative to the local band edge.

This gives the prediction:

$$\xi(S_{\text{spec}}) \approx \xi(R_1) \approx \xi(F_S) \approx \xi_s, \quad (37)$$

consistent with the experimental observation that all descriptors show $\xi \approx 1\text{--}2 \mu\text{m}$.

V. CORRELATION—LENGTH HIERARCHY

This section derives the central analytical result of the paper: the correlation-length hierarchy $\xi(E_{\text{cent}}) \geq \xi(E_{\text{dom}})$. We state it as a proposition and prove it from the decomposition of the dominant-peak energy into a slow part and a short-range trap-switching part. As emphasized in the introduction, the informative content is not the mere ordering (any average is smoother than a peak position) but its *magnitude*, which encodes the relative strength of the trap and slow disorder.

A. Main result

Proposition 3 (Correlation-Length Hierarchy). For the model (1)–(6) with $W_s, W_t > 0$:

$$\xi(E_{\text{cent}}) \geq \xi(E_{\text{dom}}) \gtrsim \xi_t \equiv \sqrt{2}\sigma_t, \quad (38)$$

where ξ_t is the trap correlation length. The first inequality follows within the present decomposition from the positivity of the short-range trap-switching variance (Sec. V); the second is a microscopic floor that should be read as an order-of-magnitude lower bound. In practice, $\xi(E_{\text{dom}})$ is not set by ξ_t itself but by the larger trap-switching scale $\ell_{\text{switch}} \sim \sigma_{\text{opt}}$ and is further broadened by optical and pixel (Δx) averaging, so the effective floor is $\gtrsim \max(\xi_t, \sigma_{\text{opt}}, \Delta x)$. The equality $\xi(E_{\text{cent}}) = \xi(E_{\text{dom}})$ holds when the trap-switching variance σ_s^2 vanishes, for example, in the limit $W_t \rightarrow 0$ (or when the dominant peak is always set by the smooth background, so that traps never switch E_{dom}).

B. Proof

Step 1: Decomposition of dominant-peak energy. The dominant-peak energy E_{dom} is set by the position of the highest spectral peak within the optical spot. We write

$$E_{\text{dom}}(\mathbf{R}) = \bar{V}_s(\mathbf{R}) + \delta E_{\text{dom}}(\mathbf{R}), \quad (39)$$

where $\delta E_{\text{dom}}(\mathbf{R})$ captures the deviation of the dominant peak from the smooth background. Unlike the centroid, which averages over all states, the dominant-peak energy is set by the single state with the highest optical weight at position $\mathbf{R}$ —which is typically a trap state or a localized moiré state near $\mathbf{R}$.

Step 2: Statistics of δE_{dom} . The fluctuations δE_{dom} arise from (1) switching of trap levels within the optical spot as $\mathbf{R}$ varies and (2) variation in which moiré-localized state has

the highest weight. Both sources produce fluctuations with correlation lengths $\ell_{\text{switch}} \sim \sigma_{\text{opt}}$ (scale over which trap configuration changes as the optical spot moves) and $\ell_t \sim \sigma_t$ (individual trap scale). Since $\sigma_t \lesssim \sigma_{\text{opt}} \ll \xi_s$:

$$\xi \delta E_{\text{dom}} \sim \sigma_{\text{opt}} \ll \xi_s. \quad (40)$$

Step 3: Correlation length of E_{dom} . The total E_{dom} correlation function is

$$\begin{aligned} \langle E_{\text{dom}}(\mathbf{R}) E_{\text{dom}}(\mathbf{R}') \rangle &= \langle \bar{V}_s(\mathbf{R}) \bar{V}_s(\mathbf{R}') \rangle \\ &+ \langle \delta E_{\text{dom}}(\mathbf{R}) \delta E_{\text{dom}}(\mathbf{R}') \rangle \\ &= W_s^2 e^{-r^2/2\xi_{\text{eff}}^2} + \sigma_\delta^2 h(r/\sigma_{\text{opt}}), \end{aligned} \quad (41)$$

where $r = |\mathbf{R} - \mathbf{R}'|$ (as in Corollary 1), $\sigma_\delta^2 = \text{Var}(\delta E_{\text{dom}})$, and $h(x) \rightarrow 0$ for $x \gg 1$. In writing Eq. (41), we have neglected the cross-correlation $\langle \bar{V}_s(\mathbf{R}) \delta E_{\text{dom}}(\mathbf{R}') \rangle$. Weak trap clustering correlates the trap positions with V_{slow} and makes this term nonzero, but it is slaved to $\bar{V}_s$ and therefore only renormalizes the effective short-range variance σ_δ^2 ; it cannot make E_{dom} smoother than E_{cent} , so the ordering $\xi(E_{\text{cent}}) > \xi(E_{\text{dom}})$ is preserved. The variance satisfies

$$\text{Var}(E_{\text{dom}}) = W_s^2 + \sigma_\delta^2 > W_s^2 = \text{Var}(E_{\text{cent}}), \quad \sigma_\delta > 0. \quad (42)$$

The normalized correlation function

$$C_{E_{\text{dom}}}(r) = \frac{W_s^2 e^{-r^2/2\xi_{\text{eff}}^2} + \sigma_\delta^2 h(r/\sigma_{\text{opt}})}{W_s^2 + \sigma_\delta^2}, \quad (43)$$

decays from 1 to 0 at a scale smaller than ξ_{eff} because the $\sigma_\delta^2 h(r/\sigma_{\text{opt}})$ term vanishes quickly while the remaining $W_s^2 e^{-r^2/2\xi_{\text{eff}}^2} / (W_s^2 + \sigma_\delta^2) < 1$ at $r = 0$.

At the $1/e$ scale of the experimentally relevant hierarchy, $\xi(E_{\text{dom}}) \gtrsim 1 \mu\text{m} \gg \sigma_{\text{opt}}$, so the short-range term $\sigma_\delta^2 h(r/\sigma_{\text{opt}})$ has already decayed and only the smooth background contribution survives. The $1/e$ length $\xi(E_{\text{dom}})$ then satisfies the implicit equation

$$\frac{W_s^2}{W_s^2 + \sigma_\delta^2} e^{-\xi(E_{\text{dom}})^2/2\xi_s^2} = e^{-1}. \quad (44)$$

Solving Eq. (44) in the perturbative regime $\sigma_\delta^2/W_s^2 < 1$ gives a shortened correlation length $\xi(E_{\text{dom}}) < \xi_s \approx \xi(E_{\text{cent}})$; in the weak-trap-noise limit $\sigma_\delta^2 \ll W_s^2$, this reduces to the closed-form expansion

$$\xi(E_{\text{dom}}) \approx \xi_s \left(1 - \frac{\sigma_\delta^2}{2W_s^2} \right). \quad (45)$$

Outside the perturbative regime, $\xi(E_{\text{dom}})$ must be obtained directly from Eq. (41); only the qualitative ordering $\xi(E_{\text{dom}}) < \xi(E_{\text{cent}})$ is preserved within the present decomposition.

Order-of-magnitude estimate. From the experiment [35]: $\xi(E_{\text{cent}}) \approx 2.00 \mu\text{m}$, $\xi(E_{\text{dom}}) \approx 1.27 \mu\text{m}$, so $\xi(E_{\text{dom}})/\xi(E_{\text{cent}}) \approx 0.64$. Inserting this ratio into the Gaussian implicit relation (44) [with $\xi(E_{\text{cent}}) \approx \xi_s$] and solving for the trap-switching variance gives

$$\frac{\sigma_\delta^2}{W_s^2} = \exp \left[1 - \frac{\xi(E_{\text{dom}})^2}{2\xi_s^2} \right] - 1 \approx 1.2, \quad (46)$$

i.e., a trap-switching fluctuation $\sigma_\delta \sim W_s$ of order unity. The independent estimate from the centroid-dominant correlation [Eq. (49) below, $\sigma_\delta/W_s \approx 0.8$] falls in the same range; the residual spread reflects the assumed functional form of the mixed correlation decay, so we read this only as an order-of-magnitude estimate. Both routes place the trap disorder comparable to the slow-disorder amplitude, in the hierarchical regime, and the simulated correlation functions [Fig. 3(a)] confirm the hierarchy directly, as summarized by the correlation-length bar chart [Fig. 3(b)].

Range of validity. The weak-noise expansion [Eq. (45)] is strictly controlled only for $\sigma_\delta^2/W_s^2 \ll 1$. The experimentally inferred ratio $\sigma_\delta^2/W_s^2 \approx 1.2$ lies outside this perturbative regime, so the numerical inversion above should be read as an *order-of-magnitude estimate* based on the full two-component correlation form Eq. (41), not as a small-noise expansion. The sign of the hierarchy $\xi(E_{\text{dom}}) < \xi(E_{\text{cent}})$ is robust within the present decomposition for any $\sigma_\delta > 0$, following from the positivity of the trap-switching variance; only the precise mapping $\sigma_\delta^2/W_s^2 \leftrightarrow \xi(E_{\text{dom}})/\xi_s$ is approximation controlled.

VI. INTERDESCRIPTOR CORRELATION THEORY

By an *interdescriptor correlation*, we mean the correlation between two spatial descriptor maps $X(\mathbf{R})$ and $Y(\mathbf{R})$ evaluated over all measured pixels $\mathbf{R}$ of the hyperspectral map; it thus measures how two coarse descriptors of the local spectrum covary across the sample, not a correlation between individual spectral peaks. Throughout this section, correlations between two such descriptor fields $X(\mathbf{R})$ and $Y(\mathbf{R})$ are quantified by the Pearson coefficient

$$\rho_P(X, Y) = \frac{\text{Cov}(X, Y)}{\sqrt{\text{Var}(X) \text{Var}(Y)}}, \quad (47)$$

where $\text{Cov}(X, Y) = \langle XY \rangle - \langle X \rangle \langle Y \rangle$ and $\text{Var}(X) = \text{Cov}(X, X)$, with $\langle \cdot \rangle$ denoting the spatial average over the map. The experimental descriptor matrix, in contrast, is reported using the Spearman rank-correlation coefficient ρ_S [35], defined as the Pearson correlation of the rank-transformed descriptor fields,

$$\rho_S(X, Y) = \rho_P(\text{rank } X, \text{rank } Y), \quad (48)$$

where $\text{rank } X$ denotes the rank-transformed values of $X(\mathbf{R})$ over all map pixels. It thus measures monotonic association without assuming a linear relation. The Spearman coefficient is the natural choice here because the spectral descriptors are nonlinear functionals of the local PL spectrum and need not be Gaussian distributed; rank correlations therefore provide a robust measure of monotonic covariation that is directly comparable to the experimental descriptor-correlation matrix. We compute ρ_P analytically because it is exact for the Gaussian disorder variables underlying our model and use it as a sign-and-magnitude estimate for ρ_S : For the monotonic descriptor relations considered here, the two coefficients share the same sign and have comparable magnitude, so that the analytical ρ_P predicts the measured ρ_S . Where we quote our own simulation values for comparison with experiment (e.g., Sec. XID), we report ρ_S to match the experimental convention.




FIG. 3. Spatial correlation-length hierarchy confirmed by simulation. (a) Normalized spatial correlation functions $C(r)$ for six spectral descriptors, computed from a single 20×20 hyperspectral map on an $8 \times 8 \mu\text{m}^2$ domain with best-fit parameters ($W_s = 6 \text{ meV}$, $\xi_s = 2 \mu\text{m}$, $W_t = 12 \text{ meV}$). The horizontal dotted line marks the $1/e$ threshold, and the legend lists the $1/e$ length of each highlighted descriptor for this single realization [the ensemble means over 30 realizations are shown in panel (b)]. Thin light-gray lines are the remaining, unhighlighted descriptors (ΔE_{cd} , R_{HL} , I_{tot}). The dark-gray solid curve is the whole-spectrum cosine similarity. Because every pixel shares the same broad envelope, its raw value sits on a high nonzero plateau; we therefore subtract that long-range plateau and renormalize to unity at $r = 0$, so that only the physically decaying excess is shown. Its $1/e$ length, $1.85 \mu\text{m}$, matches the experimental whole-spectrum value of $1.87 \mu\text{m}$. (b) Simulated (blue) vs experimental (red) $1/e$ correlation lengths. The simulated bars are the mean $\pm$ standard deviation over 30 disorder realizations, with ξ extracted as the $1/e$ decay length of the spatial autocorrelation (the experimental convention). The simulation confirms the hierarchy $\xi(E_{\text{cent}}) > \xi(E_{\text{dom}})$ (blue arrow): $\xi(E_{\text{cent}}) = 1.93 \pm 0.14 \mu\text{m}$ and $\xi(E_{\text{dom}}) = 1.38 \pm 0.22 \mu\text{m}$ (ratio 0.72 ± 0.10), consistent with the experimental 2.00 and 1.27 (ratio 0.64) [35]. The F_S and R_1 lengths fall below the experimental values, reflecting the Poisson-limited trap statistics of the simplified model (Sec. XII); their finite-size robustness is examined in Appendix F.

A. Centroid-dominant correlation

From Eqs. (31) and (39), the Pearson correlation between the centroid and peak-position fields is

$$\begin{aligned} \rho_P(E_{\text{cent}}, E_{\text{dom}}) &= \frac{\text{Cov}(\bar{V}_s, \bar{V}_s + \delta E_{\text{dom}})}{\sqrt{\text{Var}(\bar{V}_s)} \sqrt{\text{Var}(\bar{V}_s) + \text{Var}(\delta E_{\text{dom}})}} \\ &= \frac{W_s}{\sqrt{W_s^2 + \sigma_\delta^2}} < 1. \end{aligned} \quad (49)$$

Here, W_s denotes the optically filtered slow-disorder amplitude $\text{Var}(\bar{V}_s)^{1/2}$ (Sec. III), which is indistinguishable from the bare W_s in the limit $\xi_s \gg \sigma_{\text{opt}}$. Identifying this with the experimental Spearman value [35] $\rho_s \approx 0.79$ gives $\sigma_\delta/W_s \approx 0.8$, of the same order as the correlation-length estimate above ($\sigma_\delta/W_s \sim 1$); both indicate $\sigma_\delta \sim W_s$.

B. The $\Delta E_{\text{cd}}\text{-}R_{\text{HL}}$ anticorrelation

The strong experimental anticorrelation $\rho_s(\Delta E_{\text{cd}}, R_{\text{HL}}) \approx -0.978$ [35] is the most striking interdescriptor correlation. We now derive this as a robust spectral shape relation for spectra with a dominant unimodal envelope.

Proposition 4 ($\Delta E_{\text{cd}}\text{-}R_{\text{HL}}$ anticorrelation). Consider spectra with a dominant unimodal envelope whose shape is governed by a single asymmetry parameter α , with $\alpha = 0$ the case symmetric about the dominant peak E_{dom} . The centroid-dominant offset $\Delta E_{\text{cd}} = E_{\text{cent}} - E_{\text{dom}}$ and the high/low ratio R_{HL} are then monotone in α with opposite sense: ΔE_{cd} and

$R_{\text{HL}} - 1$ vanish together at $\alpha = 0$ and carry opposite signs for $\alpha \neq 0$,

$$\begin{aligned} \Delta E_{\text{cd}} < 0 &\iff R_{\text{HL}} > 1, \\ \Delta E_{\text{cd}} > 0 &\iff R_{\text{HL}} < 1, \end{aligned}$$

so that $\rho_P(\Delta E_{\text{cd}}, R_{\text{HL}}) \rightarrow -1$ when α is the dominant source of pixel-to-pixel variation. (Absent the single-parameter restriction, the sign relation is a strong tendency rather than an identity, since ΔE_{cd} weights the below/above intensities by their mean energy distances whereas R_{HL} weights them equally.)

Proof. The centroid $E_{\text{cent}} = \int E I(E) dE / \int I(E) dE$ weights high-energy spectral weight more heavily relative to E_{dom} when the spectrum has a high-energy tail. A low-energy tail concentrates weight below E_{dom} , giving $R_{\text{HL}} > 1$ and pulling the centroid down so that $E_{\text{cent}} < E_{\text{dom}}$ and $\Delta E_{\text{cd}} < 0$. Conversely, a high-energy tail gives $R_{\text{HL}} < 1$ and $\Delta E_{\text{cd}} > 0$.

More precisely, define $m_1^+ = \int_{E_{\text{dom}}}^\infty E I(E) dE$ and $m_1^- = \int_{-\infty}^{E_{\text{dom}}} E I(E) dE$. Then,

$$\begin{aligned} E_{\text{cent}} - E_{\text{dom}} &= \frac{m_1^+ + m_1^-}{I_{\text{tot}}} - E_{\text{dom}} \\ &= \frac{m_1^+ - E_{\text{dom}} I_{\text{tot}}^>}{I_{\text{tot}}} + \frac{m_1^- - E_{\text{dom}} I_{\text{tot}}^<}{I_{\text{tot}}}, \end{aligned} \quad (50)$$

where $I_{\text{tot}}^{>}<$ are the integrated intensities above/below E_{dom} . Both terms represent signed “moments” of spectral weight

about E_{dom} , which are monotone functions of spectral asymmetry.

The ratio $R_{\text{HL}} = I_{\text{tot}}^{\leq}/I_{\text{tot}}^{\geq}$ is a monotone function of the same asymmetry parameter α . Since both ΔE_{cd} and $R_{\text{HL}} - 1$ are monotone in α and vanish at $\alpha = 0$ with opposite sign, their joint locus is a monotone curve of negative slope, giving $\rho_P(\Delta E_{\text{cd}}, R_{\text{HL}}) \rightarrow -1$ in the limit of pure asymmetry variation. This argument refers to the asymmetry of the broad envelope about its dominant maximum, not to the placement of individual lines: An isolated low-energy trap peak can itself set E_{dom} and produce $\Delta E_{\text{cd}} > 0$ (as in the trap-rich pixels of Fig. 6), but the observed near-one-dimensional $\Delta E_{\text{cd}}-R_{\text{HL}}$ relation indicates that it is the envelope asymmetry, rather than such discrete level switching, that remains the controlling variable across the map. Because both descriptors are referenced to the dominant peak, the anticorrelation requires this shared reference: In our simulated maps, $\rho_S(\Delta E_{\text{cd}}, R_{\text{HL}}) = -0.97 \pm 0.01$ with the E_{dom} split, but only $+0.50$ with an E_{cent} split and $+0.11$ with a fixed-median split—showing that the strong relation is tied to the dominant-peak-referenced definition (the experimentally adopted one), rather than a consequence of an arbitrary normalization choice. ■

Remark 2. The near-perfect anticorrelation $\rho_S \approx -0.978$ observed experimentally indicates that spectral asymmetry variation—rather than peak multiplicity or peak energy modulation—is the dominant source of descriptor variation in this dataset. This is a nontrivial diagnostic: The mere observation that $\rho_S(\Delta E_{\text{cd}}, R_{\text{HL}}) \approx -1$ is difficult to reconcile with simple models in which multiple incoherent peaks at unrelated energies dominate the spectrum, and instead establishes that the broad envelope remains effectively unimodal across the map—directly confirming the two-fluid picture of Sec. VIII.

C. Sharp fraction-roughness correlation

Both F_S and R_1 measure the density of sharp spectral structure. They differ in how the structure is weighted: F_S measures the fraction of integrated intensity in narrow features, while R_1 measures the total variation of the spectral profile. For a spectrum consisting of a smooth background plus sharp peaks, both increase with the number and depth of sharp peaks, giving a large positive correlation.

Analytically, for a spectrum $I = I_{\text{bg}} + \sum_i A_i g_\eta(E - E_i)$:

$$R_1 \approx R_1^{\text{bg}} + \frac{2}{\eta\sqrt{2\pi}} \times \frac{\sum_i A_i}{I_{\text{tot}}}, \quad (51)$$

$$F_S \approx \frac{\sum_i A_i}{I_{\text{tot}}}, \quad (52)$$

giving $R_1 \approx R_1^{\text{bg}} + \text{const} \times F_S$, and thus $\rho_P(F_S, R_1) \rightarrow 1$ when trap peak variation dominates. The experimental value $\rho_S \approx 0.901$ [35] confirms this.

D. Width-entropy correlation

The quantile width W_{80} and spectral entropy S_{spec} both measure the spread of spectral weight. For a spectrum spanning N_{eff} effective resolution elements, $W_{80} \propto N_{\text{eff}}$ and $S_{\text{spec}} \propto \ln N_{\text{eff}}$. Since both are monotone functions of N_{eff} , they are positively correlated. The experimental value $\rho_S \approx 0.846$ [35] is consistent. The full simulated Spearman interdescriptor




FIG. 4. Spearman interdescriptor correlation matrix from simulation. Each cell shows the Spearman rank correlation coefficient ρ_S between a pair of spectral descriptors, computed from a simulated 20×20 hyperspectral map with best-fit parameters. Gold boxes mark the four pairs also reported experimentally [35]— $\Delta E_{\text{cd}}-R_{\text{HL}}$, F_S-R_1 , $W_{80}-S_{\text{spec}}$, and $E_{\text{cent}}-E_{\text{dom}}$. The displayed coefficients are those of a single best-fit map; the five-seed mean simulated values and their measured counterparts are compared in Table II (mean absolute deviation 0.010 per pair). Note the strong positive block structure among $(I_{\text{tot}}, F_S, R_1)$ (trap-rich pixels) and the negative block coupling these to $(E_{\text{cent}}, E_{\text{dom}})$, reflecting the low-energy trap emission at potential minima.

matrix is shown in Fig. 4. The corresponding simulation-experiment comparison is summarized in Table II.

VII. DISORDER PARAMETER REGIMES

The qualitative behavior of the model is set by the two disorder strengths—the slow-disorder amplitude W_s and the effective trap disorder Γ_t —measured against the homogeneous linewidth η . Which of these exceeds η determines whether the local spectra are essentially featureless, smoothly modulated on the micron scale, densely peaked, or organized on both scales at once. In this section, we map that dependence onto four qualitatively distinct regimes, in the spirit of a phase diagram (Fig. 5), and use the map to locate the experimental

TABLE II. Spearman interdescriptor correlations: simulation vs experiment. Simulation values are means over five independent disorder realizations; the experimental values are from Ref. [35]. The total score $\mathcal{L} = 0.041$ represents a mean absolute deviation of 0.010 per pair.

Descriptor pair	Simulation (mean $\pm$ std)	Experiment
$(\Delta E_{\text{cd}}, R_{\text{HL}})$	-0.976 ± 0.009	-0.978
(F_S, R_1)	$+0.903 \pm 0.024$	$+0.901$
$(W_{80}, S_{\text{spec}})$	$+0.821 \pm 0.065$	$+0.846$
$(E_{\text{cent}}, E_{\text{dom}})$	$+0.800 \pm 0.075$	$+0.788$




FIG. 5. Disorder parameter map for the hierarchical disorder model. Four qualitatively distinct regimes are defined by the disorder strengths relative to the single-level linewidth $\eta = 1.5$ meV: Regime I (weak disorder, $W_s, W_t < \eta$): spectrally uniform emission; Regime II (slow dominated, $W_s > \eta > W_t$): smooth micron-scale spectral modulation; Regime III (trap dominated, $W_t > \eta > W_s$): dense random sharp lines without spatial organization; and Regime IV (hierarchical disorder dominated, $W_s, W_t > \eta$): coexisting micron-scale domains and dense local trap manifolds; the descriptor correlation-length hierarchy is maximally visible. The gold star marks the best-fit parameters reproducing the four principal Spearman descriptor correlations reported experimentally [35] simultaneously. Here, n_t and σ_t are held fixed, so the vertical axis W_t represents the trap-disorder strength $\Gamma_t^{1/2} = (n_t W_t^2 \sigma_t^2)^{1/2}$ of Eq. (53) up to a constant factor.

system. This matters because the descriptor correlation-length hierarchy studied here is prominent only when *both* disorder scales exceed η ; identifying the regime therefore tells us when the descriptor analysis applies. We first introduce the two control parameters, then define the regime boundaries, and finally place the MoSe₂/WSe₂ system on the resulting map.

A. Control parameters

We map the parameter space of the two disorder strengths into qualitatively distinct regimes:

W_s (slow disorder amplitude, meV),

$$\Gamma_t \equiv n_t W_t^2 \sigma_t^2 \quad (\text{effective trap disorder, meV}^2). \quad (53)$$

The quantity Γ_t is the effective trap-potential variance, combining the trap density n_t , the depth variance W_t^2 , and the trap area $\sim \sigma_t^2$ through the dimensionless coverage $n_t \sigma_t^2$; its square root $\Gamma_t^{1/2}$ is the rms trap-potential amplitude (in meV), directly comparable to W_s and η .

B. Regime boundaries

We define four regimes based on the hierarchy of spectral disorder. The boundaries below are order-of-magnitude

criteria intended to demarcate qualitatively distinct regimes; the exact boundaries depend on microscopic details and should be regarded as schematic. The geometric factors $(a_M/\xi_s)^{1/2}$ that appear below are schematic coarse-graining factors estimating how nanoscale trap fluctuations (on the moiré scale a_M) survive averaging over a slow-disorder domain of size ξ_s ; they are heuristic, not the result of a microscopic derivation.

(1) *Regime I: Uniform excitonic (small W_s , small Γ_t)*. Neither disorder scale produces significant spectral structure. The PL is a single smooth Gaussian-like peak at the unperturbed exciton energy. Boundary: $W_s < \eta$ (spectral resolution) and $\Gamma_t^{1/2} < \eta$.

(2) *Regime II: Smooth correlated-domain (large W_s , small Γ_t)*. Slow disorder creates micron-scale spectral domains. The spectrum at each pixel is smooth (low roughness, low entropy, small F_S). E_{cent} maps show smooth spatial modulation with $\xi(E_{\text{cent}}) \approx \xi_s$. Boundary: $W_s > \eta$ and $\Gamma_t^{1/2} < W_s \times (a_M/\xi_s)^{1/2}$.

(3) *Regime III: Random line-rich (small W_s , large Γ_t)*. Dense sharp spectral peaks appear at random positions. High roughness and entropy, but no spatial domain structure. Boundary: $W_s < \eta$ and $\Gamma_t^{1/2} > \eta$.

(4) *Regime IV: Hierarchical disorder-dominated (large W_s , large Γ_t)*. Both disorder scales are active. Micron-scale spectral domains coexist with local multipeak manifolds. The descriptor correlation-length hierarchy $\xi(E_{\text{cent}}) > \xi(E_{\text{dom}})$ is maximally observable and directly measurable in this regime:

Regime IV: $W_s > \eta$,

$$\Gamma_t^{1/2} > W_s \times (a_M/\xi_s)^{1/2}. \quad (54)$$

C. Location of the experimental system

The experimental data directly constrain the model parameters. The measured $\xi(E_{\text{cent}}) \approx 2.00$ μm implies $\xi_s \approx 2$ μm ; the centroid energy variation across the map implies $W_s \sim 10$ – 15 meV; and the prevalence of multipeak spectra with $F_S > 0.1$ and $R_1 \sim 0.3$ – 0.5 meV⁻¹ confirms $\Gamma_t^{1/2} > \eta \sim 2$ meV. The experimental system therefore falls within Regime IV (hierarchical disorder dominated), as illustrated in the disorder parameter map (Fig. 5). The value $W_s \sim 10$ – 15 meV here is the *raw centroid variation* across the experimental map. In the numerical simulations of Sec. XI, we use $W_s = 6$ meV as an *effective* slow-disorder amplitude that accounts for optical-spot averaging and finite-map normalization, so that the two values refer to different operational definitions of the slow-disorder strength rather than a quantitative inconsistency.

VIII. TWO-FLUID LDOS STRUCTURE

The eigenstates of $H = -(\hbar^2/2M)\nabla^2 + V_{\text{slow}} + V_{\text{trap}}$ fall into two populations based on their energy relative to the local slow potential.

Background states $[E_n \lesssim \bar{V}_s(\mathbf{R}_n)]$ are weakly bound to slow-disorder minima, with spatial extents $\ell_n \gg \sigma_t$. Each contributes a broad Lorentzian of width η_{bg} centered near $\bar{V}_s(\mathbf{R}_n)$. The incoherent sum of ~ 20 background eigenstates

per optical spot produces a smooth, quasi-Gaussian envelope that tracks $\bar{V}_s(\mathbf{R})$ and constitutes the broad background I_{bg} .

Trap states [those lying more than $\sim W_t$ below the local background, $E_n \lesssim \bar{V}_s(\mathbf{R}_i) - W_t$] are spatially concentrated near individual trap sites $\mathbf{R}_i$, with spatial extent $\ell_n \sim \sigma_t$. Each contributes a narrow peak of width $\eta \ll \eta_{\text{bg}}$ at energy $E_n \approx \bar{V}_s(\mathbf{R}_i) + E_n^{(\text{trap})}$.

The resulting local density of states at position $\mathbf{R}$ is

$$\rho(E, \mathbf{R}) \approx \underbrace{\rho_{\text{bg}}(E - \bar{V}_s(\mathbf{R}))}_{\text{smooth background}} + \underbrace{\sum_{i: \mathbf{R}_i \in \text{spot}(\mathbf{R})} \rho_i^{\text{trap}}(E - \bar{V}_s(\mathbf{R}_i))}_{\text{sharp trap peaks}}, \quad (55)$$

and the local PL spectrum follows from integrating $\rho(E, \mathbf{R})$ against the optical weight function. This two-fluid structure is reproduced at the model level by the Hamiltonian diagonalization in Fig. 8(b): The filled areas show the background and trap contributions separately.

A key nontrivial consequence is that the broad-background intensity $I_{\text{bg,tot}}$ and sharp-peak intensity $I_{\text{sharp,tot}}$ are *positively* correlated:

$$\text{Cov}(I_{\text{bg,tot}}, I_{\text{sharp,tot}}) > 0. \quad (56)$$

This follows because both populations are enhanced in slow-potential minima: Background eigenstates are preferentially concentrated in slow-potential minima by disorder selection, while traps are preferentially nucleated at the same sites. The positive covariance distinguishes the hierarchical disorder landscape from a two-phase competition model, in which sharp peaks would form at the expense of the background.

IX. SPECTRAL FAMILIES AS CORRELATED DISORDER DOMAINS

A. Origin of spectral families

The experimental observation of three dominant spectral families in PCA + Gaussian mixture clustering raises the question: Are these families distinct thermodynamic phases or emergent statistical features of the disorder landscape?

In our framework, spectral families emerge naturally from the correlated slow-disorder field $V_{\text{slow}}(\mathbf{r})$. A Gaussian random field on a finite domain $[0, L]^2$ with correlation length $\xi_s \ll L$ develops $N_{\text{dom}} \sim (L/\xi_s)^2$ quasi-independent regions with characteristic local values $V_{\text{slow}} \sim W_s$.

For the experiment [35], $L \approx 8 \mu\text{m}$, $\xi_s \approx 2 \mu\text{m}$, giving $N_{\text{dom}} \sim 16$ statistically independent regions. However, PCA + clustering of a nine-dimensional descriptor space may resolve only the dominant modes of variation, yielding ~ 3 effective families.

Remark 3 [Spectral family number (heuristic estimate)].

We expect the number of spectral families resolved by Gaussian mixture clustering on the first k principal components of the descriptor space to be $N_{\text{fam}} \approx \min(k+1, N_{\text{dom}})$ in the regime where interdomain descriptor separation exceeds intradomain variance. The $k+1$ form is a heuristic counting estimate (each retained mode adds at most one resolvable cluster boundary), not a sharp theorem.

B. Statistical nature of spectral families

The spectral families in the present framework are therefore *not* sharply distinct thermodynamic phases (like phase-separated domains with a definite interfacial free energy). Instead, they are emergent statistical clusters of the continuously varying slow disorder landscape, identified by a finite number of PCA modes and separated by soft boundaries that represent continuous crossovers rather than sharp interfaces.

This interpretation carries several testable consequences. The cluster boundaries should be smooth, and the number of identifiable families should grow with the PCA truncation order rather than saturating at a thermodynamically fixed count. Different spectral descriptors should produce consistent but not identical family assignments, since each descriptor filters a different combination of disorder scales. The characteristic domain size of each family should equal $\sim \xi_s$, in agreement with the experimentally observed mean domain diameter of $1.79 \mu\text{m}$ [35].

X. MULTISCALE SPECTRAL ORGANIZATION IN THE HIERARCHICAL REGIME

The hierarchically disordered regime (Regime IV, both $W_s \gg \eta$ and $\Gamma_t \gg \eta^2$) exhibits multiscale spectral organization that is absent in any single-disorder model. The effective potential $V_{\text{eff}} = V_{\text{slow}} + V_{\text{trap}}$ supports a nested landscape: The smooth background set by V_{slow} organizes the micron-scale spectral envelope, while V_{trap} provides the fine structure within each optical-spot region.

A quantitative diagnostic for the degree of multiscale organization is the mean spectral cosine similarity between spatially separated pixels. For the minimal model, the spatial correlation of the full spectra $C_l(\mathbf{r}) \equiv \langle I(E, \mathbf{R}) \cdot I(E, \mathbf{R} + \mathbf{r}) \rangle / \langle \|I\|^2 \rangle$ decays on a length scale set by ξ_s . The baseline-subtracted cosine similarity of our simulated map (Fig. 3) decays with a $1/e$ length of $1.85 \mu\text{m}$, in close agreement with the whole-spectrum cosine correlation length of $1.87 \mu\text{m}$ reported experimentally [35]. The same landscape that produces the descriptor-specific correlation lengths (Sec. V) also governs the domain structure visible in PCA cluster maps (Sec. IX).

The key observable consequence is that the spectral covariance is not self-averaging on the experimental map: Different $8 \times 8 \mu\text{m}^2$ realizations produce statistically similar descriptor distributions but different spatial patterns, reflecting the finite ratio $L/\xi_s \approx 4$ of map size to correlation length. This non-self-averaging is a direct signature of the multiscale organization and distinguishes the hierarchical regime from both the smooth-domain (Regime II) and trap-dominated (Regime III) limits.

XI. NUMERICAL SIMULATIONS

A. Overview: Two complementary approaches

Numerical validation of the analytical predictions proceeds via two complementary routes on a common 128×128 grid spanning an $8 \mu\text{m} \times 8 \mu\text{m}$ domain ($\Delta x = 62.5 \text{ nm}$).

Route 1 (primary): Hamiltonian diagonalization. The Hamiltonian acts as a coarse-grained *spectral generator*: It

defines a spatially correlated eigenvalue spectrum whose density of states, when convolved with the optical point-spread function, produces a well-posed LDOS at every map position. The key property is that the spectral descriptors at a given position depend on which eigenstates are in competition for spectral weight within the optical spot—a many-eigenstate quantity that requires solving the full eigenvalue problem rather than reading off local potential values.

Concretely, the full sparse exciton Hamiltonian $H = -(\hbar^2/2M)\nabla^2 + V_{\text{slow}} + V_{\text{trap}}$ is assembled and diagonalized for its lowest optically active eigenstates via the Lanczos algorithm (ARPACK) [64]; for larger systems, the same low-lying spectrum can be obtained more economically with shift-and-invert sparse eigensolvers such as SLEPC/PETSc [65,66], allowing a finer spatial grid. The diagonalization yields eigenstates (ε_n, ψ_n) from which the local PL spectrum is constructed as $I(E, \mathbf{R}_j) = \sum_n A_{jn} g_{\sigma_n}(E - \varepsilon_n)$, with the incoherent local-density optical weights A_{jn} [Eq. (16)]; Appendix E confirms that the coherent weight (13) yields statistically indistinguishable descriptor statistics within numerical uncertainty. This route verifies that the correlation-length hierarchy and interdescriptor Spearman relations *emerge from eigenvalue statistics* without relying on the synthetic two-fluid spectrum construction of Route 2. In the limit of vanishing hopping (kinetic) energy $t_{\text{hop}} \equiv \hbar^2/(2M\Delta x^2) \rightarrow 0$ (the finite-difference hopping of Appendix B), the model reduces to a spatially correlated random spectral manifold with no quantum transport; the finite t_{hop} acts solely as a spectral regularizer that converts the discrete disorder landscape into a continuous, normalizable LDOS. Results are presented in Sec. XIF.

Route 2 (computational): Phenomenological PL model. As a computationally efficient alternative enabling large-ensemble parameter sweeps, we directly synthesize local PL spectra from the two-fluid physical picture as

$$I(E, \mathbf{R}) = I_{\text{bg}}(E; \bar{V}_s(\mathbf{R})) + f_{\text{trap}} \sum_{k \in \mathcal{N}(\mathbf{R})} w_k(\mathbf{R}) g_{\eta}(E - E_k), \quad (57)$$

where $\bar{V}_s(\mathbf{R}) = \int W_{\text{opt}}(\mathbf{r} - \mathbf{R}) V_{\text{slow}}(\mathbf{r}) d^2\mathbf{r}$ is the optically averaged slow potential, I_{bg} is a Gaussian of width $\eta_{\text{bg}} = 10\eta$, $w_k(\mathbf{R}) = \exp(-|\mathbf{R}_k - \mathbf{R}|^2/2\sigma_{\text{opt}}^2)$ is the optical weight at trap k , $E_k = \bar{V}_s(\mathbf{R}_k) + u_k$ is the trap energy, and $f_{\text{trap}} = 0.2$ is chosen so that $F_S \sim 0.15$ – 0.40 . The model parameters f_{trap} and η_{bg} are not independently derived from the Hamiltonian; they are calibrated to the experimental sharp-fraction range. We emphasize that f_{trap} is not used to fit the descriptor correlation signs or magnitudes: It only sets the overall sharp-fraction scale $\langle F_S \rangle$ and enters the inverse parameter extraction only through that calibration, while the main correlation signs and relative magnitudes are controlled primarily by the disorder parameters $\{W_s, \xi_s, W_t, n_t, E_{\text{bias}}\}$. This route enables multiseed parameter sweeps for the Spearman-correlation comparison (Sec. XID); the correlation-length statistics in Fig. 3 and Appendix F are evaluated over 30 independent disorder realizations.

B. Trap clustering model

If trap positions are drawn from an uncorrelated Poisson process, then $\xi(F_S) \sim \sigma_{\text{opt}} \ll \xi_s$. Experimentally, however,

$\xi(F_S) \approx \xi(E_{\text{cent}}) \approx \xi_s$, suggesting that the trap density tracks the slow disorder landscape. This is physically motivated by the possibility that reconstruction domains, stacking-registry variations, strain concentrators, or charged defects correlate the locations, or the optical activation, of traplike emitters with the slow exciton-energy landscape.

We model this by the *minimal correlated-trap ansatz*: Trap positions are drawn from

$$P(\mathbf{r}) \propto \exp\left(-\frac{V_{\text{slow}}(\mathbf{r})}{E_{\text{bias}}}\right), \quad (58)$$

where $E_{\text{bias}} > 0$ is a trap-slow-disorder bias energy scale (not a temperature) that sets the degree of spatial correlation between optically active traps and slow-potential minima. This is the lowest-order ansatz that captures the coupling between the two disorder scales without introducing additional clustering parameters beyond E_{bias} ; the microscopic processes behind the correlation (defect migration, strain coupling) are not modeled explicitly. In the limit $E_{\text{bias}} \rightarrow \infty$, the distribution is uniform (random traps). For $E_{\text{bias}} \sim W_s$, over typical one-sigma variations of V_{slow} the trap density changes by a factor of order $e^{W_s/E_{\text{bias}}}$; rare $\pm 3\sigma$ excursions ($\Delta V \sim 3W_s$) can produce a much larger contrast, $n(\mathbf{r})/\langle n \rangle \sim e^{\pm 3}$. This creates domains of high and low trap density with correlation length ξ_s .

The spatial correlation of the trap-density field is

$$C_n(r) \equiv \text{Corr}[n(\mathbf{R}), n(\mathbf{R} + \mathbf{r})] = \frac{\exp[(W_s/E_{\text{bias}})^2 C_s(r)] - 1}{\exp[(W_s/E_{\text{bias}})^2] - 1}, \quad (59)$$

which inherits a correlation scale of order ξ_s for finite bias strength. This expression follows from the moment-generating function of the Gaussian field V_{slow} applied to the log-normal trap-density field $n(\mathbf{r}) \propto e^{-V_{\text{slow}}(\mathbf{r})/E_{\text{bias}}}$. For jointly Gaussian $V_{\text{slow}}(\mathbf{R})$, $V_{\text{slow}}(\mathbf{R} + \mathbf{r})$ with covariance $W_s^2 C_s(r)$, one has $\langle n(\mathbf{R})n(\mathbf{R} + \mathbf{r}) \rangle / \langle n \rangle^2 = \exp[(W_s/E_{\text{bias}})^2 C_s(r)]$, which normalizes to the correlation coefficient above. The observable F_S inherits this correlation length when the clustering signal exceeds the Poisson noise in trap count per optical spot.

C. Parameter optimization

The model has five parameters: $\{W_s, \xi_s, W_t, n_t, E_{\text{bias}}\}$ (with $\sigma_t = 50$ nm and $\eta = 1.5$ meV fixed). We determine these by minimizing the total deviation from the four experimentally reported Spearman interdescriptor correlations:

$$\mathcal{L}(\theta) = \sum_{(a,b)} |\rho_{\text{sim}}(a, b) - \rho_{\text{exp}}(a, b)|, \quad (60)$$

where the sum runs over the four pairs: $(\Delta E_{\text{cd}}, R_{\text{HL}})$, (F_S, R_1) , $(W_{80}, S_{\text{spec}})$, and $(E_{\text{cent}}, E_{\text{dom}})$. This objective directly tests the theoretical predictions of Sec. VI.

The best-fit parameters (averaged over five independent disorder realizations) are

$$W_s = 6.0 \text{ meV}, \quad \xi_s = 2.0 \mu\text{m}, \quad W_t = 12.0 \text{ meV}, \\ n_t = 2.0 \mu\text{m}^{-2}, \quad E_{\text{bias}} = 7.0 \text{ meV}. \quad (61)$$

The ratio $W_s/W_t = 0.5$ places the system in the trap-dominated energy-spread regime (Regime IV of Sec. VII),




FIG. 6. Simulated local PL spectra and spectral descriptor maps. Upper row: Three representative PL spectra from the 20×20 measurement map, illustrating the three spectral types present in the hierarchical disorder regime. (a) Background-dominated pixel: A single broad Gaussian background peak with $E_{\text{cent}} \approx E_{\text{dom}}$ and no resolved sharp lines ($F_S \approx 0$). (b) Mixed pixel: The background peak is visible alongside one or two sharp trap peaks; the centroid E_{cent} is pulled below the dominant peak E_{dom} by the trap-related weight on the low-energy side of the envelope. (c) Trap-rich pixel: Several sharp lines above the background, with the centroid pulled toward the cluster of trap energies. Vertical dashed and dotted lines mark E_{cent} (red) and E_{dom} (blue) for each pixel. (d) Scatter plot of ΔE_{cd} vs R_{HL} for all 400 pixels, colored by F_S . The tight anticorrelation ($\rho_S = -0.955$ for this single representative seed; cf. the five-seed mean $\rho_S = -0.976 \pm 0.009$ in Table II) reflects the robust spectral shape relation derived in Proposition 4; pixels with larger F_S (yellow) show the strongest separation between centroid and dominant-peak energy. Lower row: Spatial maps of four spectral descriptors across the 20×20 measurement grid. (e) $\delta E_{\text{cent}} = E_{\text{cent}} - \langle E_{\text{cent}} \rangle$: Smooth micron-scale fluctuations tracking V_{slow} , with $\xi(E_{\text{cent}}) \approx 1.93 \mu\text{m}$. (f) $\delta E_{\text{dom}} = E_{\text{dom}} - \langle E_{\text{dom}} \rangle$: More fragmented pattern with additional trap-switching fluctuations, confirming $\xi(E_{\text{dom}}) < \xi(E_{\text{cent}})$. Panels (e) and (f) share a common color scale (in meV), so the larger fluctuation amplitude of δE_{dom} is apparent. (g) F_S map: Trap-density-modulated sharp-fraction distribution. (h) R_1 (spectral roughness) map: Closely tracks F_S , consistent with $\rho_S(F_S, R_1) = 0.909$.

consistent with the experimental observation of dense multi-peak spectra. The slow-disorder amplitude $W_s = 6.0$ meV sets the centroid energy variation scale, while the trap depth $W_t = 12.0$ meV determines the dominant-energy variation and the $E_{\text{cent}}-E_{\text{dom}}$ decorrelation ratio.

The low-trap-fraction limit ($f_{\text{trap}} = 0.2 \ll 1$) predicts $\rho_P(E_{\text{cent}}, E_{\text{dom}}) \rightarrow W_s / \sqrt{W_s^2 + \bar{W}_t^2}$, where $\bar{W}_t^2 = W_t^2(1 - 2/\pi) \approx 0.363W_t^2$ for the half-normal distribution. With $W_s = 6$, $W_t = 12$: theoretical $\rho_{P,\text{min}} = 6/\sqrt{36 + 52.3} = 0.638$, which rises to ≈ 0.79 at finite $f_{\text{trap}} = 0.2$ through the partial $E_{\text{dom}}-E_{\text{bg}}$ alignment, matching the observed value.

D. Key numerical results

Figure 6 summarizes the representative PL spectra and spectral descriptor maps from the best-fit simulation: three representative pixel spectra [Figs. 6(a)–6(c)], the $\Delta E_{\text{cd}}-R_{\text{HL}}$ scatter plot [Fig. 6(d)], and the spatial maps of δE_{cent} , δE_{dom} , F_S , and R_1 [Figs. 6(e)–6(h)].

With the parameters in Eq. (61), the five-seed mean Spearman correlations are (the Spearman sweep uses five seeds for the parameter scan, whereas the correlation-length statistics in Fig. 3 and Appendix F use 30 independent disorder realizations).

Within the phenomenological PL model, all four correlations are reproduced to within the seed-to-seed standard deviation, confirming the theoretical predictions of Sec. VI. The $\Delta E_{\text{cd}}-R_{\text{HL}}$ anticorrelation is reproduced to within 0.002 ($\rho_{\text{sim}} = -0.976 \pm 0.009$ versus $\rho_{\text{exp}} = -0.978$), consistent with the robust spectral shape relation of Proposition 4. The F_S-R_1 correlation is equally well matched ($\rho_{\text{sim}} = +0.903 \pm 0.024$ versus $+0.901$); the small seed-to-seed variability reflects the discrete nature of trap counting within each optical spot. The $W_{80}-S_{\text{spec}}$ correlation is slightly underestimated ($\rho_{\text{sim}} = +0.821 \pm 0.065$ versus $+0.846$), a discrepancy attributable to the absence of the moiré potential, which would contribute additional spectral structure and broaden W_{80} . Finally, the $E_{\text{cent}}-E_{\text{dom}}$ correlation, the most sensitive diagnostic of the trap-to-slow disorder ratio, is reproduced within

TABLE III. Comparison of single-scale and hierarchical disorder models. Numerical values are computed from the same 20×20 hyperspectral simulation pipeline with identical W_s , W_t , ξ_s , n_t , and η as in the rest of this section, varying only which components are switched on. Experimental values are from Ref. [35].

Model	$\xi(E_{\text{cent}})$ (μm)	$\xi(E_{\text{dom}})$ (μm)	$\rho_S(\Delta E_{cd}, R_{\text{HL}})$	$\langle F_S \rangle$
Slow only ($W_t = 0$)	1.98	1.98	-0.47	0.00
Trap only ($W_s = 0$) ^b	0.23	0.18	-0.98	0.34
Hierarchical	1.93	1.38	-0.97	0.30
Experiment ^a	2.00	1.27	-0.978	0.15–0.40

^aRef. [35].

^bCorrelation lengths are the $1/e$ decay length of the spatial autocorrelation (matching the experimental convention), averaged over 30 disorder realizations. For the trap-only model, the descriptor fields are short range and near degenerate, so their autocorrelation does not reach $1/e$ within the map; the quoted trap-only lengths use a Gaussian fit $C(r) = e^{-r^2/2\xi^2}$, which resolves the short trap scale.

the seed-to-seed uncertainty ($\rho_{\text{sim}} = +0.800 \pm 0.075$ versus $+0.788$).

The simulated 20×20 hyperspectral map also confirms the correlation-length hierarchy (Proposition 3): $\xi(E_{\text{cent}}) = 1.93 \pm 0.14 \mu\text{m} > \xi(E_{\text{dom}}) = 1.38 \pm 0.22 \mu\text{m}$ (mean $\pm$ standard deviation over 30 disorder realizations, using the $1/e$ decay length of the spatial autocorrelation—estimated from the empirical semivariogram [67]—to match the experimental convention), with ratio $\xi(E_{\text{dom}})/\xi(E_{\text{cent}}) = 0.72 \pm 0.10$ (experiment: $1.27/2.00 = 0.64$ [35]). The corresponding standard errors of the ensemble means are approximately $0.03 \mu\text{m}$ for E_{cent} and $0.04 \mu\text{m}$ for E_{dom} , much smaller than the $\approx 0.55 \mu\text{m}$ separation between the two mean correlation lengths, so the ordering is well resolved by the ensemble. The centroid correlation length now agrees with the experimental $2.00 \mu\text{m}$; the finite-size robustness of the hierarchy—and of the ratio—is examined in Appendix F.

The simulated correlation lengths $\xi(F_S) = 0.84 \pm 0.23 \mu\text{m}$ and $\xi(R_1) = 0.93 \pm 0.20 \mu\text{m}$ remain below the experimental values of $\sim 2 \mu\text{m}$. This discrepancy is a known limitation of the simplified random-trap model (see Sec. XII), where with $n_t = 2 \mu\text{m}^{-2}$ there are only ~ 2 – 3 traps per optical spot, causing Poisson noise to partially mask the spatial trap-density correlation. The real moiré system contains $\sim 10^3$ – 10^4 moiré-scale sites per optical spot, and even a small optically active fraction may be spatially organized by reconstruction, strain, or stacking-registry variations. Such correlations provide a natural route toward $\xi(F_S) \sim \xi_s$, beyond the Poisson-limited random-trap model.

Details of the numerical implementation are provided in the analysis scripts available upon reasonable request (see the Data Availability statement).

E. Necessity of hierarchical disorder: Single-scale comparison

A natural question is whether either disorder component alone could reproduce the observed phenomenology. To address this, we repeat the phenomenological simulation in three settings, keeping all other parameters fixed at the best-fit values: (1) *slow only* ($W_t = 0$, no traps); (2) *trap only* ($W_s = 0$, no slow background); and (3) *hierarchical* (both, as above). The results are summarized in Table III and Fig. 7, which contrast the correlation functions of the slow-only [Fig. 7(a)], trap-only [Fig. 7(b)], and hierarchical [Fig. 7(c)] models.

The *slow-only* model fails because E_{cent} and E_{dom} become statistically identical: With no traps, the spectrum at every pixel is a single broad envelope centered on $\bar{V}_s(\mathbf{r})$, so $\rho_S(E_{\text{cent}}, E_{\text{dom}}) = +1.00$ and $\xi(E_{\text{cent}}) = \xi(E_{\text{dom}}) = 1.98 \mu\text{m}$. No correlation-length hierarchy is present, and $\langle F_S \rangle = 0$ (no sharp peaks).

The *trap-only* model fails in the opposite way: In the absence of a smooth background, E_{cent} and E_{dom} both inherit only the spatial scale of the optical spot, giving $\xi(E_{\text{cent}}) = 0.23 \mu\text{m}$ and $\xi(E_{\text{dom}}) = 0.18 \mu\text{m}$, both well below the




FIG. 7. Single-scale disorder cannot reproduce the correlation-length hierarchy. Spatial correlation functions $C(r)$ of E_{cent} (blue circles) and E_{dom} (red squares) for the three disorder models, shown for one representative realization; the corresponding ensemble $1/e$ correlation lengths are listed in Table III. (a) Slow only: The two descriptors are statistically degenerate; their correlation functions are indistinguishable, giving no hierarchy. (b) Trap only: Both decay within the optical spot ($\xi < \sigma_{\text{opt}}$); no micron-scale descriptor domains form. (c) Hierarchical disorder: $\xi(E_{\text{cent}}) > \xi(E_{\text{dom}})$, reproducing the experimentally observed ordering. Horizontal dotted line marks $C(r) = 1/e$.

experimental $\sim 1.27\text{--}2.00\ \mu\text{m}$ [35] and below the optical spot size σ_{opt} . Pixels develop sharp lines ($\langle F_S \rangle = 0.34$) and a strong $\rho_S(\Delta E_{\text{cd}}, R_{\text{HL}}) = -0.98$, but no micron-scale descriptor domains appear, in contradiction with the experiment.

Only the *hierarchical* model simultaneously reproduces (1) the correlation-length ordering $\xi(E_{\text{cent}}) > \xi(E_{\text{dom}})$ at the micron scale, (2) the strong $\Delta E_{\text{cd}}\text{--}R_{\text{HL}}$ anticorrelation, and (3) the observed sharp-fraction values. This is the most direct statement of the need for two disorder scales: Descriptor covariance alone argues against either single-scale model from a hyperspectral PL map.

F. Continuum LDOS validation of the descriptor hierarchy

To confirm that the correlation-length hierarchy and the principal Spearman relations emerge from *eigenvalue statistics* rather than from the construction of the phenomenological PL model, we repeated the analysis using the full sparse exciton Hamiltonian (1) diagonalized numerically.

(1) *Setup*. We build the kinetic-energy operator as a finite-difference Laplacian (Appendix B) on the 128×128 grid and add $V_{\text{eff}} = V_{\text{slow}} + V_{\text{trap}}$ on the diagonal, obtaining a $16\,384 \times 16\,384$ sparse matrix. The 1500 lowest eigenstates (ε_n, ψ_n) are computed with the Lanczos algorithm (ARPACK) [64] using the same disorder realization and best-fit parameters (61).

With $M = 1.2m_0$ and $\Delta x = 62.5\ \text{nm}$, the hopping energy is $t_{\text{hop}} = \hbar^2/(2M\Delta x^2) = 0.0081\ \text{meV}$, 4 orders of magnitude smaller than $W_s = 6\ \text{meV}$ and $W_t = 12\ \text{meV}$. The system is therefore in the disorder-dominated spectral-statistics regime ($t_{\text{hop}}/W_t \approx 7 \times 10^{-4}$): The disorder potential far exceeds the kinetic energy at the grid scale, so the eigenstate energy statistics are determined primarily by the disorder landscape. The Hamiltonian therefore serves as a coarse-grained *spectral generator*: The kinetic term defines the continuum LDOS manifold, while the disorder landscape controls the spatial arrangement of eigenstates at scales relevant to the optical spot. Solving the eigenvalue problem is then essential because the spectral descriptors at a given position are determined by the competition among many eigenstates whose weight falls inside the optical spot—a many-eigenstate quantity that cannot be read off from local potential values. Because the IPR is dominated by disorder rather than quantum interference in this regime, we use an *energy-based classification* instead: State n is traplike if $\varepsilon_n < V_{\text{slow}}(\mathbf{R}_n) - 0.7W_t$ (i.e., lying more than $8.4\ \text{meV}$ below the local background), where $\mathbf{R}_n = \arg \max_{\mathbf{r}} |\psi_n(\mathbf{r})|^2$ is the localization center. Of the 1500 lowest states, 126 are classified as traplike, close to the 133 trap positions in the disorder realization. Background states receive linewidth $\eta_{\text{bg}} = 15\ \text{meV}$; trap states receive $\eta = 1.5\ \text{meV}$ (same values as the phenomenological model).

(2) *Results*. The optical-weight matrix $A_{jn} = \sum_{\mathbf{r}} W_{\text{opt}}(\mathbf{r} - \mathbf{R}_j) |\psi_n(\mathbf{r})|^2$ is computed by a single matrix multiply, and PL spectra are synthesized as

$$I(E, \mathbf{R}_j) = \sum_n A_{jn} g_{\sigma_n}(E - \varepsilon_n), \quad (62)$$

where $\sigma_n = \eta/\sqrt{8 \ln 2}$ or $\sigma_n = \eta_{\text{bg}}/\sqrt{8 \ln 2}$ depending on state type. The resulting Spearman correlations

are

$$\begin{aligned} \rho_S(\Delta E_{\text{cd}}, R_{\text{HL}}) &= -0.971 \quad [\text{exp: } -0.978], \\ \rho_S(F_S, R_1) &= +0.941 \quad [\text{exp: } +0.901], \\ \rho_S(E_{\text{cent}}, E_{\text{dom}}) &= +0.496 \quad [\text{exp: } +0.788]. \end{aligned}$$

The first two correlations agree with experiment to within 0.04, confirming that the $\Delta E_{\text{cd}}\text{--}R_{\text{HL}}$ anticorrelation and the $F_S\text{--}R_1$ covariation emerge from eigenvalue statistics without relying on the synthetic two-fluid spectrum construction. The correlation-length hierarchy is also reproduced: $\xi(E_{\text{cent}}) = 1.18\ \mu\text{m} > \xi(E_{\text{dom}}) = 0.69\ \mu\text{m}$, with ratio of 0.58 (experiment: 0.64).

The $\rho_S(E_{\text{cent}}, E_{\text{dom}})$ value is underestimated (+0.496 versus +0.788). This is a structural consequence of the trap-dominated spectral regime: The narrow trap-state peaks ($\sigma = 0.64\ \text{meV}$) are $\sim 10\times$ taller than the broad background peaks ($\sigma = 6.4\ \text{meV}$) for the same optical weight, so E_{dom} is almost always set by the deepest trap in the optical spot rather than by the smooth background. The phenomenological model avoids this artefact by assigning the background peak amplitude independently of the trap amplitude (via $f_{\text{trap}} = 0.2$). Physically, the discrepancy reflects the fact that in the real material the radiative linewidths of individual eigenstates depend on their spatial character (phonon-bath coupling, oscillator strength) in ways that go beyond the present disorder-statistics framework.

Taken together, the two routes bracket the real system: The Hamiltonian diagonalization operates in the disorder-dominated limit (eigenstates determined purely by the potential landscape, uniform radiative coupling), while the phenomenological model operates in the effective optical-weight limit (trap intensities assigned independently of eigenstate character). The experimental system lies between these limits, and the fact that both routes reproduce the $\Delta E_{\text{cd}}\text{--}R_{\text{HL}}$ and $F_S\text{--}R_1$ spectral-shape correlations to within 0.04 confirms that these correlations are robust features of the disorder hierarchy rather than artefacts of either modeling choice (the $E_{\text{cent}}\text{--}E_{\text{dom}}$ value, discussed above, is the one correlation that is regime sensitive). For the same reason, the two routes are not expected to yield identical absolute correlation lengths [phenomenological: $\xi(E_{\text{cent}}) = 1.93 \pm 0.14\ \mu\text{m}$, $\xi(E_{\text{dom}}) = 1.38 \pm 0.22\ \mu\text{m}$; diagonalization: $\xi(E_{\text{cent}}) = 1.18\ \mu\text{m}$, $\xi(E_{\text{dom}}) = 0.69\ \mu\text{m}$]: the robust predictions are the hierarchy and the ratio, both reproduced consistently.

Figures 8 and 9 summarize the results. Figure 8 traces the emergence chain from the disorder landscape [Fig. 8(a)], through the eigenstate PL spectra [Fig. 8(b)], to the δE_{cent} [Fig. 8(c)] and δE_{dom} [Fig. 8(d)] descriptor maps, whose spatial correlation functions and $1/e$ lengths are compared in Figs. 9(a) and 9(b), respectively.

(3) *Two-fluid spectral structure at the single-pixel level*. Figure 8(b) demonstrates that the broad-background plus sharp-peak structure emerges directly from the eigenstate decomposition, without the phenomenological amplitude assignment used in Route 2. Background-dominated pixels show a broad quasi-Gaussian envelope from ~ 20 background eigenstates per optical spot, each broadened by $\eta_{\text{bg}} = 15\ \text{meV}$. Mixed and trap-rich pixels show narrow peaks




FIG. 8. Emergence chain: $H \rightarrow$ eigenstates $\rightarrow$ local spectra $\rightarrow$ descriptor maps. All panels derive from the same diagonalization of $H = -(\hbar^2/2M)\nabla^2 + V_{\text{slow}} + V_{\text{trap}}$ on a 128×128 , $8 \times 8 \mu\text{m}$ grid (best-fit parameters). (a) Disorder landscape V_{slow} with trap positions (gold $\times$). Squares, triangles, and circles mark the three representative measurement pixels. Rings show the $1/e$ radius of the optical spot (FWHM = $0.85 \mu\text{m}$). (b) PL spectra from Hamiltonian eigenstates at the three representative pixels. Filled areas show background-eigenstate contribution (shaded, broad) and trap-eigenstate contribution (solid, narrow); dashed/dotted vertical lines mark E_{cent} and E_{dom} . The trap-rich pixel shown ($F_S = 1$) is an extreme representative of the diagonalization route, chosen to illustrate the trap-dominated limit; it is not the typical sharp fraction, which averages to the experimental range $\langle F_S \rangle \approx 0.15$ – 0.40 over the full map. (c) E_{cent} spatial map ($\xi = 1.18 \mu\text{m}$): Smoothly tracks V_{slow} . (d) E_{dom} spatial map ($\xi = 0.69 \mu\text{m}$): More fragmented due to trap-level switching. The annotation confirms $\xi(E_{\text{cent}}) > \xi(E_{\text{dom}})$ emerging from eigenvalue statistics without relying on the synthetic two-fluid spectrum construction.

($\eta = 1.5 \text{ meV}$) from individual trap eigenstates rising above the broad envelope. This two-fluid coexistence is a direct consequence of the spatial-energy structure of the Hamiltonian eigenstates, not of phenomenological amplitude assignment.

XII. DISCUSSION

A. Limitations of the minimal model

The present model intentionally omits valley physics (K and K' exciton species) [14,68–70], spin structure and bright/dark exciton mixing, phonon-assisted emission and momentum-indirect transitions [71,72], exciton-exciton interactions, and nonequilibrium relaxation dynamics [73,74]—all of which are secondary to the disorder-induced spectral organization at the spatial scales of interest.

The justification for this omission is structural rather than incidental: Valley polarization, bright-dark mixing, and

phonon-assisted sidebands act *within* the single-pixel emission envelope, generating fine structure that can shift, split, or skew individual peaks. They do not, however, modify the spatial low-pass filtering of V_{slow} by the optical spot, which is what sets the macroscopic statistical moments—the centroid E_{cent} and the correlation-length hierarchy of Proposition 3. Their leading effect therefore falls on the fine-structure descriptors (F_S , R_1 , and the asymmetry encoded in ΔE_{cd} , R_{HL}), where they may renormalize the prefactors of the variance relations (C5), while leaving the spatial organization and its correlation-length hierarchy—the central results of this work—intact.

Several of these approximations can become quantitatively important in specific regimes. Valley-coherent effects may become relevant at very low trap density and temperatures below $\sim 5 \text{ K}$ [14,69], while high pump-power conditions introduce additional inhomogeneity through exciton-exciton




FIG. 9. Correlation-length hierarchy from Hamiltonian diagonalization. (a) Spatial correlation functions $C(r)$ of four spectral descriptors computed from the 400-pixel map generated by diagonalizing H . E_{cent} (blue circles) decays more slowly than E_{dom} (red squares), confirming $\xi(E_{\text{cent}}) = 1.18 \mu\text{m} > \xi(E_{\text{dom}}) = 0.69 \mu\text{m}$. The double arrow marks the hierarchy gap $\Delta\xi$. (b) Bar chart comparing diagonalization results (blue) with experimental values (red) of Ref. [35] for the four descriptors with reported correlation lengths. The hierarchy $\xi(E_{\text{cent}}) > \xi(E_{\text{dom}})$ is reproduced by both approaches, with diagonalization giving ratio of 0.58 vs experimental 0.64.

interactions. The Gaussian correlation kernel for V_{slow} also fails to capture non-Gaussian disorder such as the sharp domain walls arising from moiré reconstruction.

A more fundamental limitation is the low trap density $n_t = 2 \mu\text{m}^{-2}$, which places only ~ 2 – 3 traps within each optical spot and allows Poisson counting noise to suppress the simulated $\xi(F_s)$ and $\xi(R_1)$ (≈ 0.6 – $0.9 \mu\text{m}$) well below the experimental values ($\approx 2 \mu\text{m}$). In the real moiré system, the effective trap density is $\sim 10^3$ – 10^4 times higher, since all moiré sites contribute, and one expects $\xi(F_s) \rightarrow \xi_s$ in this high-density limit. The clustering model (58) partially mitigates this discrepancy by correlating trap positions with V_{slow} minima, but does not fully close the gap at experimentally relevant densities.

B. Non-Gaussian slow disorder

Real moiré heterobilayers often exhibit reconstruction into triangular domains with sharp boundaries [52,53,56]. In such systems, V_{slow} may be better modeled as a piecewise-constant field with sharp domain walls rather than a Gaussian random field.

The domain-wall model predicts steplike E_{cent} maps with sharp jumps between domain energies, bimodal or trimodal descriptor distributions in place of the smooth distributions expected for Gaussian disorder, and spectral family boundaries that coincide spatially with the domain walls.

This is a testable distinction: Gaussian disorder gives smooth spectral maps, while reconstruction domains give steplike maps.

A distinct, milder generalization concerns the *shape* of the correlation kernel rather than the field statistics. The Gaussian kernel of Eq. (4) is the natural choice for slow disorder arising from long-wavelength strain and electrostatic fluctuations, but

if the landscape is instead dominated by extended dislocation lines or domain-wall networks, the correlation is better described by an exponential kernel $C_s(r) = e^{-r/\xi_s}$. Importantly, the central smoothing result—that E_{cent} is an optically filtered copy of V_{slow} (Corollary 1) and that the correlation-length hierarchy $\xi(E_{\text{cent}}) \geq \xi(E_{\text{dom}})$ holds (Proposition 3)—is insensitive to this choice. Convolution with the optical spot is a low-pass filter that suppresses high spatial frequencies regardless of the kernel shape, so the qualitative filter structure is unchanged: Only the precise functional form of $C(r)$ near the origin, and hence the numerical prefactor relating $\xi(E_{\text{cent}})$ to ξ_s , depends on whether the kernel is Gaussian or exponential. The disorder-filter framework is therefore robust to the detailed form of the slow-disorder correlations.

C. Microscopic origin and magnitude of the effective parameters

The effective disorder amplitudes inferred here—a slow-disorder amplitude $W_s \sim 10$ – 15 meV and an effective trap-switching amplitude $W_t^{\text{eff}} \sim 12$ meV—are not free of microscopic constraints: They should be compatible with realistic mechanisms in a $\text{MoSe}_2/\text{WSe}_2$ heterobilayer. Although a full first-principles evaluation is beyond the scope of this effective theory, the required energy scales are consistent with reported ranges for these materials, as summarized in Table IV. Micron-scale variations of the local twist angle and long-wavelength heterostrain shift the interlayer-exciton energy through the deformation potential by several to a few tens of meV [29,54,56]; moiré reconstruction and stacking-registry variations produce stacking-dependent exciton shifts of 10–50 meV [52,53,56]; and dielectric and electrostatic inhomogeneity contributes an additional few to tens of meV [12,72]; hyperspectral PL imaging has recently begun to

TABLE IV. Characteristic energy scales in MoSe₂/WSe₂ heterobilayers that contribute to the effective disorder amplitudes. The inferred slow-disorder amplitude $W_s \sim 10\text{--}15\text{ meV}$ and effective trap-switching amplitude $W_t^{\text{eff}} \sim 12\text{ meV}$ lie within these known ranges.

Mechanism	Energy scale	Contributes to
Local twist-angle variation $\delta\theta$	Few–10 meV	V_{slow}
Heterostrain $\varepsilon \sim 0.1\%\text{--}1\%$	Few–tens meV	V_{slow}
Reconstruction/stacking registry	10–50 meV	$V_{\text{slow}}, \text{traps}$
Dielectric/electrostatic disorder	Few–tens meV	V_{slow}
Localized defect/strain traps	5–50 meV	V_{trap}

map these strain and disorder fields directly [36]. Superposed over a micron correlation length, these combine to a slow-disorder amplitude of order 10 meV, consistent with the inferred W_s . The sharp emitters—moiré-trapped excitons, atomic defects, and strain-localized sites—carry binding energies of 5–50 meV [16,26,45], so the effective trap-switching amplitude $W_t^{\text{eff}} \sim 12\text{ meV}$ can be accounted for by a moderate portion of the reported localized-state energy scale. The inferred parameters are therefore consistent with known microscopic energy scales.

The bias energy E_{bias} that colocalizes traps with the slow-potential minima [Eq. (58)] is an energy scale (not a temperature) that sets how strongly the trap density is biased toward low-energy regions of V_{slow} . Its meaning is fixed by its two limits: As $E_{\text{bias}} \rightarrow \infty$, the biasing factor $\exp[-V_{\text{slow}}/E_{\text{bias}}]$ becomes flat and the traps reduce to an uncorrelated Poisson distribution (no clustering), while for $E_{\text{bias}} \ll W_s$ the traps are pulled almost entirely into the deepest minima of V_{slow} (extreme clustering). The best-fit $E_{\text{bias}} \approx 7\text{ meV}$ is comparable to W_s and thus lies between these extremes: An energy change of order E_{bias} alters the local trap density by a factor $\sim e$, so the fit describes a moderate bias in which defect and strain-trap formation is enhanced—but not exclusively confined—at reconstruction domains and strain concentrators. This correlation energy is of the same order as the strain, reconstruction, dielectric, and electrostatic scales collected in Table IV, and so is physically reasonable for a MoSe₂/WSe₂ heterobilayer.

D. Toward quantitative parameter extraction

A key practical strength of the descriptor-filter framework is that it provides a *peak-decomposition-free* route to infer effective disorder parameters from hyperspectral PL data, without microscopic line assignment. The procedure maps four observable statistics to four model parameters sequentially. The slow-disorder correlation length ξ_s follows directly from Corollary 1: $\xi_s \approx \xi(E_{\text{cent}}) = 2.00\text{ }\mu\text{m}$ for the experimental system. The slow-disorder amplitude is then given by $W_s \approx \text{Var}(E_{\text{cent}})^{1/2}$ [Eq. (31)], yielding $W_s \sim 10\text{--}15\text{ meV}$ from the observed centroid energy variation across the map. The relative trap strength follows from the $E_{\text{cent}}\text{--}E_{\text{dom}}$ correlation via Eq. (49), using the experimental Spearman value ρ_S as an estimate of the analytical Pearson coefficient ρ_P . The quantity the descriptor covariance determines directly is the effective trap-switching fluctuation σ_δ , which is the

dominant-energy fluctuation induced by the traps rather than the bare trap-depth standard deviation; denoting this effective trap amplitude $W_t^{\text{eff}} \equiv \sigma_\delta$ gives

$$\frac{\sigma_\delta}{W_s} = \sqrt{\frac{1}{\rho_P(E_{\text{cent}}, E_{\text{dom}})^2} - 1} \equiv \frac{W_t^{\text{eff}}}{W_s}, \quad (63)$$

yielding $\sigma_\delta/W_s \approx 0.8$ from the experimental $\rho_S = 0.788$; the independent estimate from the correlation-length ratio $\xi(E_{\text{dom}})/\xi(E_{\text{cent}}) = 0.64$ gives $\sigma_\delta/W_s \approx 1.1$ ($\sigma_\delta^2/W_s^2 \approx 1.2$) via the Gaussian relation [Eq. (44)]. Both place σ_δ at order W_s , consistent with an effective $W_t^{\text{eff}} \gtrsim W_s$ (experimental values from Ref. [35]). The four-parameter extraction quotes this W_t^{eff} as the effective trap amplitude W_t . Finally, the effective sharp-emitter density n_t is constrained by the mean sharp fraction $\langle F_S \rangle \approx 0.15\text{--}0.40$, which grows with $n_t W_t^2 \sigma_t^2$; calibrating to $\langle F_S \rangle \approx 0.25$ at best-fit $W_t = 12\text{ meV}$ and $\sigma_t = 50\text{ nm}$ gives $n_t \approx 2\text{ }\mu\text{m}^{-2}$. Note that n_t is an effective parameter representing the density of optically resolvable sharp-emitter sites under the experimental conditions, not a direct count of microscopic defects (which would be orders of magnitude higher in a moiré lattice). This effective density also carries an excitation- and temperature dependence: Under higher pump power, a larger fraction of trap and moiré sites becomes optically populated, so n_t^{eff} grows with the exciton filling factor and ultimately saturates as state filling exhausts the available sites—bounded from above by the $\sim 10^3\text{--}10^4$ moiré sites contained within each optical spot. At low temperature, conversely, rapid relaxation funnels excitons into the deepest traps and reduces the number of *distinct* resolvable sharp emitters. The value $n_t \approx 2\text{ }\mu\text{m}^{-2}$ should therefore be read as the effective sharp-emitter density specific to the excitation power and temperature of the experiment, rather than as a fixed material constant.

The four extracted parameters (ξ_s, W_s, W_t, n_t) specify the leading disorder sector of the model and can be cross-validated by checking whether the predicted descriptor correlation length ratio and Spearman matrix reproduce the observed values—an internal consistency check that does not require any additional input. The remaining quantities entering the simulation—the trap radius σ_t , the broadenings η and η_{bg} , and the trap-clustering parameter E_{bias} —are not part of this minimal four-parameter extraction: σ_t enters only through the combination $n_t W_t^2 \sigma_t^2$ (Sec. VII) and is fixed at a representative suboptical value, while E_{bias} controls the spatial colocalization of traps with V_{slow} and is separately constrained by the correlation lengths of F_S and R_1 . This parameter extraction protocol should be applicable to TMD heterobilayer systems imaged by hyperspectral PL mapping [1,2,25], providing a spectroscopic disorder diagnostic independent of microscopic peak identification.

E. Generic versus moiré-specific content

It is important to delineate which of our results are generic to any multi-scale-disordered emitter and which are specific to a moiré heterobilayer. The descriptor-filter *mechanism* itself is general: The correlation-length hierarchy $\xi(E_{\text{cent}}) \geq \xi(E_{\text{dom}})$, the $\Delta E_{\text{cd}}\text{--}R_{\text{HL}}$ anticorrelation, and the descriptor covariance structure follow from the coexistence

of a smooth, micron-correlated background and a dense set of sharp, subspot emitters, and would arise equally in a monolayer TMD or a disordered quantum well with the same two-scale spectral content. Indeed, the underlying disorder-localized COM-exciton picture and the optically filtered correlation functions are essentially those developed for quantum-well excitons [61–63]; in that sense, the framework is a spatially resolved, descriptor-based counterpart of the mobility-edge phenomenology known there. The energy-dependent localization it implies has recently been observed directly in the same material system as a localization-to-delocalization transition of moiré excitons [75,76], and localized interlayer excitons occur in MoSe₂/WSe₂ even in the absence of a moiré potential [77]—underscoring that the localization physics underlying our descriptors is generic rather than moiré specific.

What the moiré heterobilayer provides is a natural and especially rich realization of the required two-scale structure. First, the moiré superlattice can provide a dense manifold of localized interlayer-exciton states and a high areal density of optically active sites—of order 10^3 – 10^4 within a single optical spot—so that the sharp-emitter ensemble needed for F_S , R_1 , and the trap-switching fluctuation σ_δ need not rely on sparse extrinsic defects. Second, moiré reconstruction, stacking-registry variation, and strain can correlate the slow potential with the locations or activation of traplike emitters, providing a plausible microscopic route to the trap- V_{slow} clustering [Eq. (58)] that lengthens $\xi(F_S)$ and $\xi(R_1)$ and can contribute to the observed spectral-family structure. At the spatial scales analyzed here, the moiré potential thus enters principally by setting the density and clustering of the local excitonic manifold (and by renormalizing the effective mass and local DOS), rather than by imprinting its own ~ 10 – 20 nm periodicity on the descriptors, which is averaged over by the optical spot. More specifically, moiré-related signatures—as opposed to generic multiscale disorder—would include the triangular reconstruction-domain organization of the descriptor maps and recurrent peaks at the moiré miniband spacing, both accessible only below the optical-spot scale (near-field or low-temperature single-site spectroscopy), as discussed next.

F. Extensions: Moiré potential and quantum emitter regime

When $V_{\text{moiré}}$ is included, the model predicts additional structure. The moiré potential introduces a period- a_M modulation of the local DOS, which may manifest as recurrent spectral peaks at energies separated by the moiré miniband spacing. When V_m exceeds the relevant thermal and linewidth scales, individual moiré sites can host localized quantum-emitter-like exciton states [30,32,46,47,78], and their spatial brightness distribution would then be modulated by the slow-disorder field V_{slow} —a testable prediction accessible via near-field or variable-temperature PL mapping.

XIII. CONCLUSION

We have developed a minimal theory for the spatial organization of spectral descriptors in moiré exciton photoluminescence, using the descriptor-based disorder-filter picture as an organizing principle for hyperspectral PL data.

Starting from an effective exciton Hamiltonian and the chain $H \rightarrow \rho(E, \mathbf{R}) \rightarrow I(E, \mathbf{R}) \rightarrow$ descriptor fields, the paper establishes the following results.

The descriptor-based disorder-filter principle provides the unifying framework: Each spectral descriptor is preferentially sensitive to a different component of the multiscale disorder landscape. The centroid energy integrates over the optical spot, filtering out subspot trap fluctuations, so that it primarily tracks the slow disorder landscape. The dominant-peak energy, in contrast, is additionally randomized by discrete trap-level switching within the optical spot. This differential filtering provides a microscopic mechanism for the observed interdescriptor differences and the correlation-length hierarchy.

The correlation-length hierarchy $\xi(E_{\text{cent}}) \geq \xi(E_{\text{dom}})$ is established analytically from the decomposition $E_{\text{dom}} = \bar{V}_s + \delta E_{\text{dom}}$. The experimental ratio $\xi(E_{\text{dom}})/\xi(E_{\text{cent}}) \approx 0.64$ [35] implies an effective trap-switching fluctuation $\sigma_\delta/W_s \approx 1.1$ (of order unity), and Hamiltonian diagonalization provides an independent model-level check of the hierarchy from the eigenvalue statistics: $\xi(E_{\text{cent}}) = 1.18 \text{ } \mu\text{m} > \xi(E_{\text{dom}}) = 0.69 \text{ } \mu\text{m}$.

The sign and approximate magnitude of the principal Spearman interdescriptor correlations are derived analytically. The near-perfect $\Delta E_{\text{cd}}\text{-}R_{\text{HL}}$ anticorrelation (measured $\rho_S = -0.978$ [35]) is a robust spectral shape relation that follows from the geometry of a broad class of spectra with a dominant unimodal envelope, while the positive $F_S\text{-}R_1$ correlation (measured $\rho_S = +0.901$ [35]) reflects both descriptors measuring the same trap-peak density. The calibrated phenomenological PL simulations reproduce all four principal correlations simultaneously, with a mean absolute deviation of 0.010 per pair.

Spectral families are interpreted as correlated disorder domains of the Gaussian slow-disorder field, rather than as thermodynamically distinct phases, with the predicted family count consistent with the experimentally observed three families [35].

Four qualitatively distinct disorder regimes are identified on a two-dimensional disorder parameter map. The experimental system falls in the hierarchically disordered regime where both slow and trap disorder exceed the homogeneous linewidth, producing multiscale spectral organization reflected in the descriptor correlation-length hierarchy.

The framework provides a peak-decomposition-free route to constrain the leading effective disorder parameters (ξ_s, W_s, W_t, n_t) from hyperspectral PL data via the descriptor covariance structure, without any microscopic line assignment (Table V). The predicted correlation-length hierarchy and interdescriptor Spearman matrix constitute directly testable signatures that can distinguish disorder regimes and guide future experiments with higher spatial resolution or variable temperature. Although motivated by MoSe₂/WSe₂, the framework should be applicable to hyperspectral PL maps in which spectra contain a smooth envelope together with unresolved local fine structure; descriptor covariance then serves as a general spectroscopic disorder diagnostic for two-dimensional materials, quantum-dot arrays, and disordered semiconductor heterostructures. Future work should address valley physics, moiré reconstruction domain walls,

TABLE V. Disorder parameter extraction from experimental observables. All quantities are derived from the descriptor covariance structure without microscopic line assignment. Here, W_t is inferred as an *effective trap-switching amplitude* $W_t^{\text{eff}} \equiv \sigma_\delta$ (the dominant-energy fluctuation induced by the traps), not as a microscopic trap-depth distribution width. The last row requires calibration within the two-component (background plus trap) PL model, because $\langle F_S \rangle$ depends on the broadenings and the relative oscillator strengths in addition to the trap density (Sec. XI D).

Observable	Extracted parameter	Formula	Value
$\xi(E_{\text{cent}})$	ξ_s	Direct	$\approx 2.00 \mu\text{m}$
$\text{Var}(E_{\text{cent}})^{1/2}$	W_s	Equation (31)	$\sim 10\text{--}15 \text{ meV}$
$\rho_S(E_{\text{cent}}, E_{\text{dom}})$	σ_δ/W_s (eff W_t/W_s)	Equation (63)	≈ 0.8
$\langle F_S \rangle$	$n_t W_t^2 \sigma_t^2$	Calibration	$n_t \approx 2 \mu\text{m}^{-2}$

and nonequilibrium relaxation in modifying the spectral landscape.

ACKNOWLEDGMENTS

We thank R. Kitaura for stimulating discussions on hyper-spectral PL organization in moiré heterobilayers. This work was supported by JSPS KAKENHI (Grants No. JP25K01609, No. JP22H05473, and No. JP21H01019), JST CREST (Grant No. JPMJCR19T1). K.W. acknowledges the financial support for Basic Science Research Projects (Grant No. 2401203) from the Sumitomo Foundation.

DATA AVAILABILITY

There are no publicly available research data or software supporting this manuscript. Requests for further information or data should be sent to the authors.

APPENDIX A: GAUSSIAN RANDOM FIELD GENERATION

The slow potential V_{slow} is a Gaussian random field with zero mean and a prescribed two-point correlation $\langle V(\mathbf{r})V(\mathbf{r}') \rangle = W^2 C(|\mathbf{r} - \mathbf{r}'|)$. We generate it with the spectral (Fourier-filtering) method, which is exact and costs only $O(N \ln N)$. The method rests on the Wiener-Khinchin theorem: For a stationary field, the power spectral density is the Fourier transform of the autocorrelation, $S(\mathbf{k}) = \int d^2\mathbf{r} C(r) e^{-i\mathbf{k}\cdot\mathbf{r}}$. Coloring spectrally flat (white) noise with $\sqrt{S(\mathbf{k})}$ therefore imprints precisely the target correlation $C(r)$.

- (1) Generate a white-noise field $\xi(\mathbf{r})$ with $\langle \xi(\mathbf{r})\xi(\mathbf{r}') \rangle = \delta^{(2)}(\mathbf{r} - \mathbf{r}')$.
- (2) Compute its Fourier transform $\tilde{\xi}(\mathbf{k})$.
- (3) Multiply by the square root of the power spectrum, $\tilde{V}(\mathbf{k}) = W\sqrt{S(\mathbf{k})}\tilde{\xi}(\mathbf{k})$.
- (4) Inverse Fourier transform: $V(\mathbf{r}) = \text{IFFT}[\tilde{V}(\mathbf{k})]$.

For the Gaussian kernel $C(r) = e^{-r^2/2\xi_s^2}$ used throughout, the power spectrum is itself Gaussian, $S(\mathbf{k}) = 2\pi\xi_s^2 e^{-\xi_s^2 k^2/2}$, so the generated field is smooth on the scale ξ_s .

Two implementation points are worth noting. First, the discrete transform carries grid-dependent prefactors (factors of $N_x N_y$ and the cell area $d\mathbf{x}^2$); rather than tracking these, we

rescale the output field to its exact target standard deviation, $V \rightarrow W V/\text{std}(V)$, which fixes the amplitude W unambiguously. Second, the FFT imposes periodic boundary conditions; because the map satisfies $L \gg \xi_s$, the resulting wrap-around correlations are negligible.

APPENDIX B: FINITE-DIFFERENCE HAMILTONIAN

We discretize the exciton Hamiltonian $H = -(\hbar^2/2M)\nabla^2 + V_{\text{eff}}$ on a 2D lattice with grid spacing a and $N_x \times N_y$ sites. The Laplacian is approximated by the standard second-order five-point stencil,

$$-\frac{\hbar^2}{2M}\nabla^2\psi(i, j) \approx \frac{\hbar^2}{2Ma^2}[4\psi(i, j) - \psi(i+1, j) - \psi(i-1, j) - \psi(i, j+1) - \psi(i, j-1)]. \quad (\text{B1})$$

Assembling this together with the potential, which is diagonal in the site basis, gives a sparse Hamiltonian: Each site (i, j) carries a diagonal entry $4t_{\text{hop}} + V_{\text{eff}}(i, j)$ and is coupled to its four nearest neighbors by a hopping $-t_{\text{hop}}$, where $t_{\text{hop}} = \hbar^2/(2Ma^2)$ is the kinetic (hopping) energy scale. The matrix has dimension $N_x N_y$ with at most five nonzero entries per row. We impose periodic boundary conditions, consistent with the periodic slow-disorder field of Appendix A.

For the simulation parameters $M = 1.2m_0$ and $\Delta x = 62.5 \text{ nm}$ (a 128×128 lattice on an $8 \mu\text{m}$ domain), $t_{\text{hop}} = \hbar^2/(2M\Delta x^2) \approx 0.0081 \text{ meV}$, consistent with Sec. II. Because $t_{\text{hop}} \ll W_s \sim 10 \text{ meV}$, the disorder energy dominates the kinetic energy at the grid scale; the kinetic term therefore plays the role of the *coarse-grained spectral generator* described in Sec. XI F rather than governing eigenstate localization.

Only the lowest optically active eigenstates are required. We obtain them with the Lanczos algorithm (ARPACK, smallest-algebraic mode) [64], as described in Sec. XI; for larger grids, a shift-and-invert sparse eigensolver (SLEPc/PETSc) [65,66] targets the same low-lying band more economically.

APPENDIX C: SPECTRAL DESCRIPTOR VARIANCE FROM DISORDER THEORY

For completeness, we collect the leading-order descriptor variances derived in the main text. They split into two groups, according to which part of the disorder each descriptor measures.

The energy-position descriptors follow from the decomposition $E_{\text{dom}} = \bar{V}_s + \delta E_{\text{dom}}$ (Secs. IV B and VI). The centroid tracks the optically filtered slow potential, whereas the dominant-peak energy carries the additional trap-switching fluctuation σ_δ ; because both share the same $\bar{V}_s$, it cancels in their difference:

$$\text{Var}(E_{\text{cent}}) \approx W_s^2, \quad (\text{C1})$$

$$\text{Var}(E_{\text{dom}}) \approx W_s^2 + \sigma_\delta^2, \quad (\text{C2})$$

$$\text{Var}(\Delta E_{\text{cd}}) \approx \sigma_\delta^2, \quad (\text{C3})$$

$$\text{Var}(R_{\text{HL}}) \propto \sigma_\delta^2. \quad (\text{C4})$$

TABLE VI. Robustness of the diagonalization results to the trap width σ_t . $\rho_{cd} \equiv \rho_S(\Delta E_{cd}, R_{HL})$, $\rho_{FR} \equiv \rho_S(F_S, R_1)$, and $\rho_{EE} \equiv \rho_S(E_{cent}, E_{dom})$. The hierarchy $\xi(E_{cent}) > \xi(E_{dom})$ holds in every row. The compensated sweep keeps $n_t \sigma_t^2$ fixed, stabilizing the loading-dependent observables $\langle F_S \rangle$ and ρ_{EE} , and thereby confirming the $n_t W_t^2 \sigma_t^2$ degeneracy.

σ_t (nm)	n_t (μm^{-2})	ξ_{cent} (μm)	ξ_{dom} (μm)	ρ_{cd}	ρ_{FR}	ρ_{EE}
Uncompensated ($n_t = 2.0$ fixed)						
50	2.00	1.18	0.69	−0.97	+0.94	+0.51
75	2.00	1.36	1.04	−0.93	+0.96	+0.70
100	2.00	1.16	0.87	−0.96	+0.94	+0.45
125	2.00	1.21	1.07	−0.98	+0.93	−0.01
Compensated ($n_t \propto \sigma_t^{-2}$)						
50	2.00	1.18	0.69	−0.97	+0.94	+0.51
75	0.89	0.98	0.71	−0.96	+0.92	+0.62
100	0.50	1.15	1.03	−0.98	+0.85	+0.82
125	0.32	1.18	1.11	−0.98	+0.86	+0.65

The fine-structure descriptors instead measure the trap loading within the optical spot. They scale with the combination $\Gamma_t = n_t W_t^2 \sigma_t^2$ (Sec. VII), reduced by the spot area σ_{opt}^2 (Sec. IV):

$$\text{Var}(F_S) \propto n_t W_t^2 \sigma_t^2 / \sigma_{opt}^2, \quad (\text{C5})$$

$$\text{Var}(R_1) \propto n_t W_t^2 \sigma_t^2 / (\eta^2 \sigma_{opt}^2), \quad (\text{C6})$$

$$\text{Var}(S_{spec}) \propto n_t W_t^2 \sigma_t^2 / \sigma_{opt}^2. \quad (\text{C7})$$

These relations give the scaling of the descriptor variances with the disorder parameters, providing additional testable predictions beyond the spatial correlation lengths. For the dimensionless descriptors (R_{HL} , F_S , S_{spec}) and for R_1 (of dimension inverse energy), the proportionality constants carry the residual dimensions—fixed by the spectral envelope width and the broadening η —so these last four relations express the parameter scaling rather than dimensionally complete identities.

APPENDIX D: ROBUSTNESS TO THE TRAP WIDTH σ_t

To verify that the effective trap width σ_t does not control the main results—as argued in Sec. II, where σ_t enters only through the degenerate combination $n_t W_t^2 \sigma_t^2$ and through the nonbinding hierarchy floor $\xi_t = \sqrt{2} \sigma_t$ —we repeat the full Hamiltonian diagonalization (Sec. XI) while sweeping σ_t over 50–125 nm, holding the disorder realization (slow field, trap positions, and trap depths) fixed through common random seeds. Table VI reports the outcome in two modes. In the *uncompensated* sweep (n_t fixed), increasing σ_t deepens and broadens the wells, raising the trap loading (and hence $\langle F_S \rangle$ and the trap-induced variance), yet the correlation-length hierarchy $\xi(E_{cent}) > \xi(E_{dom})$ and the two principal shape correlations $\rho_S(\Delta E_{cd}, R_{HL}) \approx -0.95$ and $\rho_S(F_S, R_1) \approx +0.93$ are unchanged. The E_{cent} - E_{dom} correlation, in contrast, is itself loading dependent and is not stabilized in this uncompensated sweep; its anomalously small value at $\sigma_t = 125$ nm ($\rho_{EE} = -0.01$) reflects overloading of the broadened trap wells and lies outside the best-fit regime, and is included only as a

Descriptor statistics: optical-weight choice




FIG. 10. Descriptor statistics are insensitive to the optical-weight model. For each descriptor statistic (labeled), the value obtained with the coherent radiative weight $\int W_{amp} \psi_n|^2$ (vertical axis) is plotted against that from the incoherent local-density weight $\int W_{opt} |\psi_n|^2$ (horizontal axis); error bars are the seed-to-seed standard deviation over five disorder realizations. All points lie close to the $y = x$ line, showing that the two weightings yield statistically indistinguishable descriptor covariance and correlation-length hierarchy within the seed-to-seed uncertainty.

stress test. In the *compensated* sweep, where $n_t \propto \sigma_t^{-2}$ holds the combination $n_t \sigma_t^2$ fixed, the loading-dependent quantities $\langle F_S \rangle$ and $\rho_S(E_{cent}, E_{dom})$ also stabilize, directly confirming the n_t - σ_t degeneracy. In all eight cases, the hierarchy ordering is preserved; only the magnitude of the gap narrows as σ_t grows [the trap-induced fluctuation becomes spatially broader, lengthening $\xi(E_{dom})$ toward $\xi(E_{cent})$], never inverting.

APPENDIX E: ROBUSTNESS TO THE OPTICAL-WEIGHT MODEL

The local PL weight of a COM eigenstate is, physically, the coherent radiative (oscillator-strength) matrix element $A_n(\mathbf{R}) = |\int W_{amp}(\mathbf{r} - \mathbf{R}) \psi_n(\mathbf{r}) d^2 \mathbf{r}|^2$ of Eq. (13) [61–63], rather than the incoherent local density $\int W_{opt} |\psi_n|^2$ used in the transparent LDOS form (16). Here, we verify that the descriptor statistics are insensitive to this choice. Using the same Hamiltonian eigenstates (Sec. XIV), we recompute all descriptors with both weightings for the same 128×128 , $8 \times 8 \mu\text{m}^2$ realizations, averaged over five independent disorder seeds. As shown in Fig. 10, every descriptor statistic— $\xi(E_{cent})$, $\xi(E_{dom})$, their ratio, the three principal Spearman correlations, and $\langle F_S \rangle$ —agrees between the two weightings to well within the seed-to-seed scatter [the coherent and incoherent values differ by ≤ 0.01 in every case, e.g., $\rho_S(\Delta E_{cd}, R_{HL}) = -0.928$ versus -0.929 and $\xi(E_{cent}) = 1.143$ versus $1.142 \mu\text{m}$]. The reason is structural: Both weights are localized within the optical spot, so the spatial statistics of the descriptor fields—and hence the correlation-length hierarchy and descriptor covariance that are the subject

of this work—are set by the disorder landscape, not by the detailed form of the per-state emission weight. The coherent weight does redistribute intensity toward the nodeless localized (trap) states, but this common reweighting cancels in the descriptor *correlations*.

APPENDIX F: FINITE-SIZE, FINITE-STEP, AND FINITE-SPOT ROBUSTNESS OF THE HIERARCHY

The experimental map spans $8 \times 8 \mu\text{m}^2$, only a few times the correlation lengths, so the extracted $1/e$ lengths carry finite-size uncertainty. We test the robustness of the hierarchy $\xi(E_{\text{cent}}) > \xi(E_{\text{dom}})$ by sweeping the map size L , the scan pitch, and the optical-spot width, each over 30 disorder realizations, using the fast phenomenological descriptor generator (Sec. XID) so that many configurations can be sampled. Figure 11 summarizes the domain-size sweep, showing the correlation lengths $\xi(E_{\text{cent}})$ and $\xi(E_{\text{dom}})$ versus map size [Fig. 11(a)] and the finite-size-robust hierarchy ratio [Fig. 11(b)]. Two features stand out. First, in this phenomenological sweep, the hierarchy holds across the sampled configurations, and the ratio $\xi(E_{\text{dom}})/\xi(E_{\text{cent}})$ is much less sensitive to sampling than the absolute lengths: It stays at 0.69–0.70 across $L = 8$ – $20 \mu\text{m}$ and scan pitches 0.25–0.60 μm , and varies only mildly (from 0.74 to 0.64) as the optical FWHM is swept over 0.60–1.10 μm , all bracketing the experimental ratio 0.64. Second, the *absolute* $1/e$ lengths are biased low on the smallest map and drift upward with L (a finite-size estimator bias); at the experimental map size, the simulation reproduces the measured values [$\xi(E_{\text{cent}}) = 1.92 \pm 0.17 \mu\text{m}$ versus 2.00; $\xi(E_{\text{dom}}) = 1.33 \pm 0.25 \mu\text{m}$ versus 1.27]. These differ slightly from the Fig. 3 ensemble (1.93 and 1.38 μm) because the finite-size sweep uses an independently generated set of realizations. Because finite sampling biases both lengths in the same direction, the physically meaningful ratio is comparatively stable. The principal anticorrelation is likewise robust, $\rho_S(\Delta E_{\text{cd}}, R_{\text{HL}}) = -0.97 \pm 0.01$ throughout. This supports the conclusion that the correlation-length hierarchy is not primarily a finite-sampling or coarse-gridding artifact, consistent with the ensemble-average inequality of Proposition 3; occasional




FIG. 11. Finite-size behavior of the correlation-length hierarchy. (a) $1/e$ correlation lengths $\xi(E_{\text{cent}})$ and $\xi(E_{\text{dom}})$ vs map size L (mean $\pm$ std over 30 disorder realizations; fixed 0.4 μm pitch and 0.85 μm optical FWHM); filled markers are the experimental values at $L = 8 \mu\text{m}$. (b) The hierarchy ratio $\xi(E_{\text{dom}})/\xi(E_{\text{cent}})$ is finite-size-robust (≈ 0.69 – 0.70) and close to the experimental value 0.64 (dotted line), even though the absolute lengths in panel (a) drift with L .

per-realization inversions seen in the microscopic diagonalization on a single 8 μm map are the expected finite-size scatter about this ensemble result.

- [1] E. Y. Andrei, D. K. Efetov, P. Jarillo-Herrero, A. H. MacDonald, K. F. Mak, T. Senthil, E. Tutuc, A. Yazdani, and A. F. Young, The marvels of moiré materials, *Nat. Rev. Mater.* **6**, 201 (2021).
- [2] K. F. Mak and J. Shan, Semiconductor moiré materials, *Nat. Nanotechnol.* **17**, 686 (2022).
- [3] E. C. Regan, D. Wang, E. Y. Paik, Y. Zeng, L. Zhang, J. Zhu, A. H. MacDonald, H. Deng, and F. Wang, Emerging exciton physics in transition metal dichalcogenide heterobilayers, *Nat. Rev. Mater.* **7**, 778 (2022).
- [4] A. K. Geim and I. V. Grigorieva, Van der Waals heterostructures, *Nature (London)* **499**, 419 (2013).
- [5] K. S. Novoselov, A. Mishchenko, A. Carvalho, and A. H. Castro Neto, 2D materials and van der Waals heterostructures, *Science* **353**, aac9439 (2016).
- [6] Y. Cao, V. Fatemi, S. Fang, K. Watanabe, T. Taniguchi, E. Kaxiras, and P. Jarillo-Herrero, Unconventional superconductivity in magic-angle graphene superlattices, *Nature (London)* **556**, 43 (2018).
- [7] Y. Cao, V. Fatemi, A. Demir, S. Fang, S. L. Tomarken, J. Y. Luo, J. D. Sanchez-Yamagishi, K. Watanabe, T. Taniguchi, E. Kaxiras, R. C. Ashoori, and P. Jarillo-Herrero, Correlated insulator behaviour at half-filling in magic-angle graphene superlattices, *Nature (London)* **556**, 80 (2018).
- [8] D. M. Kennes, M. Claassen, L. Xian, A. Georges, A. J. Millis, J. Hone, C. R. Dean, D. N. Basov, A. N. Pasupathy, and A. Rubio, Moiré heterostructures as a condensed-matter quantum simulator, *Nat. Phys.* **17**, 155 (2021).

- [9] K. F. Mak, C. Lee, J. Hone, J. Shan, and T. F. Heinz, Atomically thin MoS_2 : A new direct-gap semiconductor, *Phys. Rev. Lett.* **105**, 136805 (2010).
- [10] A. Splendiani, L. Sun, Y. Zhang, T. Li, J. Kim, C.-Y. Chim, G. Galli, and F. Wang, Emerging photoluminescence in monolayer MoS_2 , *Nano Lett.* **10**, 1271 (2010).
- [11] G. Eda, H. Yamaguchi, D. Voiry, T. Fujita, M. Chen, and M. Chhowalla, Photoluminescence from chemically exfoliated MoS_2 , *Nano Lett.* **11**, 5111 (2011).
- [12] G. Wang, A. Chernikov, M. M. Glazov, T. F. Heinz, X. Marie, T. Amand, and B. Urbaszek, Colloquium: Excitons in atomically thin transition metal dichalcogenides, *Rev. Mod. Phys.* **90**, 021001 (2018).
- [13] K. F. Mak and J. Shan, Photonics and optoelectronics of 2D semiconductor transition metal dichalcogenides, *Nat. Photonics* **10**, 216 (2016).
- [14] J. R. Schaibley, H. Yu, G. Clark, P. Rivera, J. S. Ross, K. L. Seyler, W. Yao, and X. Xu, Valleytronics in 2D materials, *Nat. Rev. Mater.* **1**, 16055 (2016).
- [15] H. Yu, G.-B. Liu, and W. Yao, Brightened spin-triplet interlayer excitons and optical selection rules in van der Waals heterobilayers, *2D Mater.* **5**, 035021 (2018).
- [16] K. Tran, G. Moody, F. Wu, X. Lu, J. Choi, K. Kim, A. Rai, D. A. Sanchez, J. Quan, A. Singh, *et al.*, Evidence for moiré excitons in van der Waals heterostructures, *Nature (London)* **567**, 71 (2019).
- [17] C. Jin, E. C. Regan, A. Yan, M. I. B. Utama, D. Wang, S. Zhao, Y. Qin, S. Yang, Z. Zheng, S. Shi, K. Watanabe, T. Taniguchi, S. Tongay, A. Zettl, and F. Wang, Observation of moiré excitons in WSe_2/WS_2 heterostructure superlattices, *Nature (London)* **567**, 76 (2019).
- [18] Y. Xu, S. Liu, D. A. Rhodes, K. Watanabe, T. Taniguchi, J. Hone, V. Elser, K. F. Mak, and J. Shan, Correlated insulating states at fractional fillings of moiré superlattices, *Nature (London)* **587**, 214 (2020).
- [19] Y. Tang, L. Li, T. Li, Y. Xu, S. Liu, K. Barmak, K. Watanabe, T. Taniguchi, A. H. MacDonald, J. Shan, and K. F. Mak, Simulation of Hubbard model physics in WSe_2/WS_2 moiré superlattices, *Nature (London)* **579**, 353 (2020).
- [20] Y. Shimazaki, I. Schwartz, K. Watanabe, T. Taniguchi, M. Kroner, and A. Imamoglu, Strongly correlated electrons and hybrid excitons in a moiré heterostructure, *Nature (London)* **580**, 472 (2020).
- [21] P. Rivera, J. R. Schaibley, A. M. Jones, J. S. Ross, S. Wu, G. Aivazian, P. Klement, K. Seyler, G. Clark, N. J. Ghimire, J. Yan, D. G. Mandrus, W. Yao, and X. Xu, Observation of long-lived interlayer excitons in monolayer MoSe_2 - WSe_2 heterostructures, *Nat. Commun.* **6**, 6242 (2015).
- [22] P. Rivera, K. L. Seyler, H. Yu, J. R. Schaibley, J. Yan, D. G. Mandrus, W. Yao, and X. Xu, Valley-polarized exciton dynamics in a 2D semiconductor heterostructure, *Science* **351**, 688 (2016).
- [23] P. Nagler, G. Plechinger, M. V. Ballottin, A. Mitoglu, S. Meier, N. Paradiso, C. Strunk, A. Chernikov, P. C. M. Christianen, C. Schüller, and T. Korn, Interlayer exciton dynamics in a dichalcogenide monolayer heterostructure, *2D Mater.* **4**, 025112 (2017).
- [24] M. Okada, A. Kutana, Y. Kureishi, Y. Kobayashi, Y. Saito, T. Saito, K. Watanabe, T. Taniguchi, S. Gupta, Y. Miyata, B. I. Yakobson, H. Shinohara, and R. Kitaura, Direct and indirect interlayer excitons in a van der Waals heterostructure of $\text{hBN}/\text{WS}_2/\text{MoS}_2/\text{hBN}$, *ACS Nano* **12**, 2498 (2018).
- [25] A. Ciarrocchi, F. Tagarelli, A. Avsar, and A. Kis, Excitonic devices with van der Waals heterostructures: Valleytronics meets twistronics, *Nat. Rev. Mater.* **7**, 449 (2022).
- [26] K. L. Seyler, P. Rivera, H. Yu, N. P. Wilson, E. L. Ray, D. G. Mandrus, J. Yan, W. Yao, and X. Xu, Signatures of moiré-trapped valley excitons in $\text{MoSe}_2/\text{WSe}_2$ heterobilayers, *Nature (London)* **567**, 66 (2019).
- [27] E. M. Alexeev, D. A. Ruiz-Tijerina, M. Danovich, M. J. Hamer, D. J. Terry, P. K. Nayak, S. Ahn, S. Pak, J. Lee, J. I. Sohn, M. R. Molas, M. Koperski, K. Watanabe, T. Taniguchi, K. S. Novoselov, R. V. Gorbachev, H. S. Shin, V. I. Fal'ko, and A. I. Tartakovskii, Resonantly hybridized excitons in moiré superlattices in van der Waals heterostructures, *Nature (London)* **567**, 81 (2019).
- [28] S. Brem, C. Linderälv, P. Erhart, and E. Malic, Tunable phases of moiré excitons in van der Waals heterostructures, *Nano Lett.* **20**, 8534 (2020).
- [29] J. Choi, M. Florian, A. Steinhoff, D. Erben, K. Tran, D. S. Kim, L. Sun, J. Quan, R. Claassen, S. Majumder, J. A. Hollingsworth, T. Taniguchi, K. Watanabe, K. Ueno, A. Singh, G. Moody, F. Jahnke, and X. Li, Twist angle-dependent interlayer exciton lifetimes in van der Waals heterostructures, *Phys. Rev. Lett.* **126**, 047401 (2021).
- [30] Y.-M. He, G. Clark, J. R. Schaibley, Y. He, M.-C. Chen, Y.-J. Wei, X. Ding, Q. Zhang, W. Yao, X. Xu, C.-Y. Lu, and J.-W. Pan, Single quantum emitters in monolayer semiconductors, *Nat. Nanotechnol.* **10**, 497 (2015).
- [31] M. Koperski, K. Nogajewski, A. Arora, V. Cherkez, P. Mallet, J.-Y. Veuillen, J. Marcus, P. Kossacki, and M. Potemski, Single photon emitters in exfoliated WSe_2 structures, *Nat. Nanotechnol.* **10**, 503 (2015).
- [32] A. Srivastava, M. Sidler, A. V. Allain, D. S. Lembke, A. Kis, and A. Imamoglu, Optically active quantum dots in monolayer WSe_2 , *Nat. Nanotechnol.* **10**, 491 (2015).
- [33] C. Chakraborty, L. Kinnischtzke, K. M. Goodfellow, R. Beams, and A. N. Vamivakas, Voltage-controlled quantum light from an atomically thin semiconductor, *Nat. Nanotechnol.* **10**, 507 (2015).
- [34] P. Tonndorf, R. Schmidt, R. Schneider, J. Kern, M. Buscema, G. A. Steele, A. Castellanos-Gomez, H. S. J. van der Zant, S. Michaelis de Vasconcellos, and R. Bratschitsch, Single-photon emission from localized excitons in an atomically thin semiconductor, *Optica* **2**, 347 (2015).
- [35] N. F. Ahmad, Y. Urano, K. Watanabe, T. Taniguchi, D. Kozawa, and R. Kitaura, Hierarchical spectral inhomogeneity in photoluminescence of a twisted $\text{MoSe}_2/\text{WSe}_2$ heterobilayer moiré superlattice revealed by hyperspectral mapping, *Appl. Phys. Lett.* **128**, 231901 (2026).
- [36] A. Alfrey, C. Tait, T. Taniguchi, K. Watanabe, and S. T. Cundiff, Revealing strain and disorder in transition-metal dichalcogenides using hyperspectral photoluminescence imaging, in *CLEO 2026*, Technical Digest Series (Optica Publishing Group, 2026), paper FW2A.6.
- [37] F. Wu, T. Lovorn, E. Tutuc, and A. H. MacDonald, Hubbard model physics in transition metal dichalcogenide moiré bands, *Phys. Rev. Lett.* **121**, 026402 (2018).

- [38] F. Wu, T. Lovorn, and A. H. MacDonald, Topological exciton bands in moiré heterojunctions, *Phys. Rev. Lett.* **118**, 147401 (2017).
- [39] M. H. Naik, E. C. Regan, Z. Zhang, Y.-H. Chan, Z. Li, D. Wang, Y. Yoon, C. S. Ong, W. Zhao, S. Zhao, *et al.*, Intralayer charge-transfer moiré excitons in van der Waals superlattices, *Nature (London)* **609**, 52 (2022).
- [40] D. A. Ruiz-Tijerina and V. I. Fal'ko, Interlayer hybridization and moiré superlattice minibands for electrons and excitons in heterobilayers of transition-metal dichalcogenides, *Phys. Rev. B* **99**, 125424 (2019).
- [41] M. H. Naik and M. Jain, Ultraflatbands and shear solitons in moiré patterns of twisted bilayer transition metal dichalcogenides, *Phys. Rev. Lett.* **121**, 266401 (2018).
- [42] E. C. Regan, D. Wang, C. Jin, M. I. B. Utama, B. Gao, X. Wei, S. Zhao, W. Zhao, Z. Zhang, K. Yumigeta, M. Blei, J. D. Carlström, K. Watanabe, T. Taniguchi, S. Tongay, M. Crommie, A. Zettl, and F. Wang, Mott and generalized Wigner crystal states in WSe₂/WS₂ moiré superlattices, *Nature (London)* **579**, 359 (2020).
- [43] Y. Bai, L. Zhou, J. Wang, W. Wu, L. J. McGilly, D. Halbertal, C. F. B. Lo, F. Liu, J. Ardelean, P. Rivera, N. R. Finney, X.-C. Yang, D. N. Basov, W. Yao, X. Xu, J. Hone, A. N. Pasupathy, and X.-Y. Zhu, Excitons in strain-induced one-dimensional moiré potentials at transition metal dichalcogenide heterojunctions, *Nat. Mater.* **19**, 1068 (2020).
- [44] A. Ghiotto, E.-M. Shih, G. S. S. G. Pereira, D. A. Rhodes, B. Kim, J. Zang, A. J. Millis, K. Watanabe, T. Taniguchi, J. C. Hone, L. Wang, C. R. Dean, and A. N. Pasupathy, Quantum criticality in twisted transition metal dichalcogenides, *Nature (London)* **597**, 345 (2021).
- [45] K. Parto, S. I. Azzam, K. Banerjee, and G. Moody, Defect and strain engineering of monolayer WSe₂ enables site-controlled single-photon emission up to 150 K, *Nat. Commun.* **12**, 3585 (2021).
- [46] H. Yu, M. Chen, and W. Yao, Giant magnetic field from moiré-induced Berry phase in homobilayer semiconductors, *Natl. Sci. Rev.* **7**, 12 (2020).
- [47] A. Branny, S. Kumar, R. Proux, and B. D. Gerardot, Deterministic strain-induced arrays of quantum emitters in a two-dimensional semiconductor, *Nat. Commun.* **8**, 15053 (2017).
- [48] A. Tarantola, *Inverse Problem Theory and Methods for Model Parameter Estimation* (Society for Industrial and Applied Mathematics, Philadelphia, 2005).
- [49] T. C. Berkelbach, M. S. Hybertsen, and D. R. Reichman, Theory of neutral and charged excitons in monolayer transition metal dichalcogenides, *Phys. Rev. B* **88**, 045318 (2013).
- [50] A. Chernikov, T. C. Berkelbach, H. M. Hill, A. Rigosi, Y. Li, B. Aslan, D. R. Reichman, M. S. Hybertsen, and T. F. Heinz, Exciton binding energy and nonhydrogenic Rydberg series in monolayer WS₂, *Phys. Rev. Lett.* **113**, 076802 (2014).
- [51] H. Yu, G.-B. Liu, J. Tang, X. Xu, and W. Yao, Moiré excitons: From programmable quantum emitter arrays to spin-orbit-coupled artificial lattices, *Sci. Adv.* **3**, e1701696 (2017).
- [52] A. Weston, Y. Zou, V. Enaldiev, A. Summerfield, N. Clark, V. Zolyomi, A. Graham, C. Yelgel, S. Magorrian, M. Zhou, J. Zultak, D. Hopkinson, A. Barinov, T. H. Bointon, A. Kretinin, N. R. Wilson, P. H. Beton, V. I. Fal'ko, S. J. Haigh, and R. Gorbachev, Atomic reconstruction in twisted bilayers of transition metal dichalcogenides, *Nat. Nanotechnol.* **15**, 592 (2020).
- [53] L. J. McGilly, A. Kerelsky, N. R. Finney, K. Shapovalov, E.-M. Shih, A. Ghiotto, Y. Zeng, S. L. Moore, W. Wu, Y. Bai, K. Watanabe, T. Taniguchi, M. Stengel, L. Zhou, J. Hone, X. Zhu, D. N. Basov, C. Dean, C. E. Dreyer, and A. N. Pasupathy, Visualization of moiré superlattices, *Nat. Nanotechnol.* **15**, 580 (2020).
- [54] H. J. Conley, B. Wang, J. I. Ziegler, R. F. Haglund, S. T. Pantelides, and K. I. Bolotin, Bandgap engineering of strained monolayer and bilayer MoS₂, *Nano Lett.* **13**, 3626 (2013).
- [55] A. Raja, L. Waldecker, J. Zipfel, Y. Cho, S. Brem, J. D. Ziegler, M. Kulig, T. Taniguchi, K. Watanabe, E. Malic, T. F. Heinz, T. C. Berkelbach, and A. Chernikov, Dielectric disorder in two-dimensional materials, *Nat. Nanotechnol.* **14**, 832 (2019).
- [56] S. Shabani, D. Halbertal, W. Wu, M. Chen, S. Liu, J. Hone, W. Yao, D. N. Basov, X. Zhu, and A. N. Pasupathy, Deep moiré potentials in twisted transition metal dichalcogenide bilayers, *Nat. Phys.* **17**, 720 (2021).
- [57] H.-P. Komsa and A. V. Krashenninnikov, Native defects in bulk and monolayer MoS₂ from first principles, *Phys. Rev. B* **91**, 125304 (2015).
- [58] N. R. Campbell, The study of discontinuous phenomena, *Proc. Cambridge Philos. Soc.* **15**, 117 (1909).
- [59] N. R. Campbell, Discontinuities in light emission, *Proc. Cambridge Philos. Soc.* **15**, 310 (1910).
- [60] H. Haug and S. W. Koch, *Quantum Theory of the Optical and Electronic Properties of Semiconductors*, 5th ed. (World Scientific, Singapore, 2009).
- [61] E. Runge and R. Zimmermann, Spatially resolved spectra, effective mobility edge, and level repulsion in narrow quantum wells, *Phys. Status Solidi B* **206**, 167 (1998).
- [62] E. Runge and R. Zimmermann, Level repulsion in excitonic spectra of disordered systems and local relaxation kinetics, *Ann. Phys.* **510**, 417 (1998).
- [63] R. Zimmermann, E. Runge, and V. Savona, Theory of resonant secondary emission: Rayleigh scattering versus luminescence, in *Quantum Coherence, Correlation and Decoherence in Semiconductor Nanostructures*, edited by T. Takagahara (Academic Press, San Diego, 2003), Chap. 4, pp. 89–165.
- [64] R. B. Lehoucq, D. C. Sorensen, and C. Yang, *ARPACK Users' Guide: Solution of Large-scale Eigenvalue Problems with Implicitly Restarted Arnoldi Methods* (Society for Industrial and Applied Mathematics, Philadelphia, 1998).
- [65] V. Hernandez, J. E. Roman, and V. Vidal, SLEPc: A scalable and flexible toolkit for the solution of eigenvalue problems, *ACM Trans. Math. Software* **31**, 351 (2005).
- [66] S. Balay, W. D. Gropp, L. C. McInnes, and B. F. Smith, Efficient management of parallelism in object-oriented numerical software libraries, in *Modern Software Tools for Scientific Computing*, edited by E. Arge, A. M. Bruaset, and H. P. Langtangen (Birkhäuser, Boston, 1997), pp. 163–202.
- [67] N. A. C. Cressie, *Statistics for Spatial Data* (Wiley, New York, 1993).
- [68] D. Xiao, G.-B. Liu, W. Feng, X. Xu, and W. Yao, Coupled spin and valley physics in monolayers of MoS₂ and other group-VI dichalcogenides, *Phys. Rev. Lett.* **108**, 196802 (2012).
- [69] K. F. Mak, K. He, J. Shan, and T. F. Heinz, Control of valley polarization in monolayer MoS₂ by optical helicity, *Nat. Nanotechnol.* **7**, 494 (2012).

- [70] K. Shinokita, K. Watanabe, T. Taniguchi, and K. Matsuda, Valley relaxation of the moiré excitons in a WSe₂/MoSe₂ heterobilayer, *ACS Nano* **16**, 16862 (2022).
- [71] G. Moody, C. K. Dass, K. Hao, C.-H. Chen, L.-J. Li, A. Singh, K. Tran, G. Clark, X. Xu, G. Berghäuser, E. Malic, A. Knorr, and X. Li, Intrinsic homogeneous linewidth and broadening mechanisms of excitons in monolayer transition metal dichalcogenides, *Nat. Commun.* **6**, 8315 (2015).
- [72] F. Cadiz, E. Courtade, C. Robert, G. Wang, Y. Shen, H. Cai, T. Taniguchi, K. Watanabe, H. Carrere, D. Lagarde, M. Manca, T. Amand, P. Renucci, S. Tongay, X. Marie, and B. Urbaszek, Excitonic linewidth approaching the homogeneous limit in MoS₂-based van der Waals heterostructures, *Phys. Rev. X* **7**, 021026 (2017).
- [73] M. Kulig, J. Zipfel, P. Nagler, S. Blanter, C. Schüller, T. Korn, N. Paradiso, M. M. Glazov, and A. Chernikov, Exciton diffusion and halo effects in monolayer semiconductors, *Phys. Rev. Lett.* **120**, 207401 (2018).
- [74] J. Zipfel, M. Kulig, R. Perea-Causín, S. Brem, J. D. Ziegler, R. Rosati, T. Taniguchi, K. Watanabe, M. M. Glazov, E. Malic, and A. Chernikov, Exciton diffusion in monolayer semiconductors with suppressed disorder, *Phys. Rev. B* **101**, 115430 (2020).
- [75] E. Blundo, F. Tuzi, S. Ciani, M. Cuccu, K. Olkowska-Pucko, Ł. Kipcza, G. Contestabile, A. Miriametro, M. Felici, G. Pettinari, T. Taniguchi, K. Watanabe, A. Babiński, M. R. Molas, and A. Polimeni, Localisation-to-delocalisation transition of moiré excitons in WSe₂/MoSe₂ heterostructures, *Nat. Commun.* **15**, 1057 (2024).
- [76] H. Fang, Q. Lin, Y. Zhang, J. Thompson, S. Xiao, Z. Sun, E. Malic, S. P. Dash, and W. Wiczorek, Localization and interaction of interlayer excitons in MoSe₂/WSe₂ heterobilayers, *Nat. Commun.* **14**, 6910 (2023).
- [77] F. Mahdikhany Sarvejahany, D. N. Shanks, M. Klein, Q. Wang, M. R. Koehler, D. G. Mandrus, T. Taniguchi, K. Watanabe, O. L. A. Monti, B. J. LeRoy, and J. R. Schaibley, Localized interlayer excitons in MoSe₂-WSe₂ heterostructures without a moiré potential, *Nat. Commun.* **13**, 5354 (2022).
- [78] M. Brotons-Gisbert, B. D. Gerardot, A. W. Holleitner, and U. Wurstbauer, Interlayer and moiré excitons in atomically thin double layers: From individual quantum emitters to degenerate ensembles, *MRS Bull.* **49**, 914 (2024).