The Nanoscale Structure of Charge Order in Cuprate Superconductor Bi2201

A DISSERTATION PRESENTED
BY
TATIANA A. WEBB
TO
THE DEPARTMENT OF PHYSICS

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY
IN THE SUBJECT OF
PHYSICS

HARVARD UNIVERSITY
CAMBRIDGE, MASSACHUSETTS
MAY 2019

Dissertation Advisor: Jennifer Hoffman Tatiana A. Webb

The Nanoscale Structure of Charge Order in Cuprate Superconductor Bi2201

Abstract

In 1986, high-temperature superconductivity was first discovered [1] in layered cuprate materials based on CuO_2 planes. The subsequent decades uncovered a complex phenomenology including a variety of additional unconventional electronic behaviors arising from strong electron-electron correlations. Superconductivity occurs near the meeting point of a regime characterized by mysteriously open “Fermi arcs” in place of a conventional closed Fermi surface and a metallic phase at high hole doping. Describing the doping evolution of the ground state is crucial to uncovering the mechanisms underlying the complex phase diagram. Here, we use the structure of broken translational symmetry, namely *d*-form factor charge modulations in $(Bi,Pb)_2(Sr,La)_2CuO_{6+\delta}$, as a probe of the reorganization marked by the transition from Fermi arcs to a large Fermi surface. From real-space imaging in a scanning tunneling microscope, we disentangle the spatial inhomogeneity of the electronic structure to obtain the doping dependence of both the Fermi surface and the charge modulation structure. We discover a commensurate to incommensurate transition accompanies the Fermi surface transition, reflecting the distinct nature of electronic correlations governing the two sides of this quantum phase transition.

Charge order in cuprates is one example of a physical relationship encoded in the spatial structure of inhomogeneous properties. Today's imaging technology enables collection of large spatially resolved data sets, but disentangling the complex structure of correlations among properties is challenging in the presence of strong disorder. Machine learning has revolutionized image analysis with its abilities to model large multidimensional data sets and to detect features despite strong variability. We explore the use of convolutional neural networks to uncover the relationships among multiple spatially resolved properties by training networks to classify the structure of disordered d-form factor modulations according to an underlying doping-dependence of the wave vector.

Contents

Listing of figures

Listing of tables

Citations to previously published work

Portions of the abstract and the majority of Chapter 3 will appear in Tatiana A. Webb, Michael C. Boyer, Yi Yin, Debanjan Chowdhury, Yang He, Takeshi Kondo, T. Takeuchi, H. Ikuta, Eric W. Hudson, Jennifer E. Hoffman, and Mohammad H. Hamidian, "Density Wave Probes Cuprate Quantum Phase Transition," Physical Review X **9**, 021021 (2019).

Glossary

STM data

a	CuO_2 lattice constant.	-
V	Sample bias voltage.	Section 2.1
I	Tunneling current.	Equation 2.1
$N(\mathbf{r}, E)$	Local density of states of the sample.	Section 2.1
$g(E = eV) = dI/dV$	Differential conductance. Proportional to local density of states.	Equation 2.2
$Z(E) = g(+E)/g(-E)$	Z map. Removes the setup effect. Enhances Bogoliubov quasiparticle interference.	Section 2.1
$\mathcal{R} = \left[\sum_{E>0} g(E)\right] / \left[\sum_{E<0} g(E)\right]$		Equation 2.5

$\mathcal{R}$ map. Removes the setup effect. Enhances density wave signal.

Form factors

l	Section 2.3
Generic form factor index. Takes values of s, s' or d.	
$\rho_l(\mathbf{r}) = a_l(\mathbf{r})\psi(\mathbf{r})$	Equation 2.6
Generic l-form factor modulation.	
$a_l(\mathbf{r})$	Section 2.3
Intra-unit cell structure for an l-form factor modulation.	
$\mathcal{A}_l$	Section 2.3
Amplitude of l-form factor modulation.	
φ_l	Section 2.3
Constant phase offset of l-form factor modulation.	
$M_{\mathrm{Cu}}(\mathbf{r})$, $M_{\mathrm{O}_x}(\mathbf{r})$, $M_{\mathrm{O}_y}(\mathbf{r})$	Equation 2.7
Sublattice masks.	
$S(\mathbf{r})$	Equation 2.8
s-form factor component of $\mathcal{R}(\mathbf{r})$.	
$S'(\mathbf{r})$	Equation 2.8
s'-form factor component of $\mathcal{R}(\mathbf{r})$.	
$D(\mathbf{r})$	Equation 2.8

d-form factor component of $\mathcal{R}(\mathbf{r})$.

Density wave structure

φ Chapter 2

Used to denote a constant phase.

ϕ Chapter 2

Used to denote the spatially varying phase of a modulation.

Φ Chapter 2

Used to denote the complex phase of a demodulated wave.

$\psi(\mathbf{r}) = A(\mathbf{r})\cos[\phi(\mathbf{r})]$ Equation 2.9

Generic disordered modulation.

$\Psi_{\mathbf{Q}}(\mathbf{r}) = A_{\mathbf{Q}}(\mathbf{r})\exp[i\Phi_{\mathbf{Q}}(\mathbf{r})]$ Equation 2.10

Demodulated wave. $\tilde{\Psi}_{\mathbf{Q}}(\mathbf{q}) = \exp\left[-\frac{q^2}{2\Lambda^2}\right]\tilde{\psi}\left(\mathbf{q}+\mathbf{Q}\right)$.

Λ Equation 2.10

q-space bandwidth of $\Psi_{\mathbf{Q}}(\mathbf{r})$.

$A_{\boldsymbol{Q}}(\mathbf{r})$ Equation 2.11

Amplitude of $\Psi_{\boldsymbol{Q}}(\mathbf{r})$.

$\Phi_{\mathbf{Q}}(\mathbf{r})$ Equation 2.11

Complex phase of $\Psi_{\boldsymbol{Q}}(\mathbf{r})$.

$R_{\mathbf{Q}}[\psi]$ Equation 2.12

Demodulation residue. Minimized to determine Q_{DW}.

$\bar{Q}_{\mathrm{DW}}$ Chapter 3

Sample average Q_{DW}.

α Chapters 2, 4

Generic spatial index. Takes values x or y.

$\mathbf{Q}_\alpha$ Section 4.3

Sample average wavevector ($\bar{Q}_{\mathrm{DW}}$) for α-direction modulations.

$\tilde{D}_\alpha(\mathbf{q}) = \exp\left[-\frac{|\mathbf{q}-\mathbf{Q}_\alpha|^2}{2\Lambda^2}\right]\tilde{D}(\mathbf{q})$ Equation 4.4

α-direction d-form factor density wave modulations isolated from D. $D_\alpha(\mathbf{r})$ is the inverse Fourier transform.

$F(\mathbf{r}) = \frac{|\Psi_{\mathbf{Q}_x}(\mathbf{r})|-|\Psi_{\mathbf{Q}_y}(\mathbf{r})|}{|\Psi_{\mathbf{Q}_x}(\mathbf{r})|+|\Psi_{\mathbf{Q}_y}(\mathbf{r})|}$ Equation 4.5

x to y asymmetry of d-form factor modulations in $D(\mathbf{r})$.

Gap masking

$\Delta(\mathbf{r})$ Figure 3.1

Local spectral gap. Proxy for local doping.

Δ^* Figure 3.6

Value of Δ where antinodal quasiparticle interference vanishes.

$M_b(\mathbf{r})$ Section 3.4

Mask for gap bin b.

$M_{a,b}(\mathbf{r}) = \sum_{j=a}^{b} M_j(\mathbf{r})$ Section 3.4

Additive gap mask for gap bins a–b.

$\mathcal{M}_b(\mathbf{r})$ Section 3.4

Smoothed mask for bin b.

w Section 3.4

Mask smoothing bandwidth.

$D_{a,b}(\mathbf{r}) = \mathcal{M}_{a,b}(\mathbf{r})D(\mathbf{r})$ Section 3.4

Mask applied to $D(\mathbf{r})$.

$D_b^{fw}(\mathbf{q})$, $D_b^{bw}(\mathbf{q})$ Equation 3.1

Gap-masked $D(\mathbf{r})$ (forward and backward direction masking).

$Z_b^{fw}(\mathbf{q})$, $Z_b^{bw}(\mathbf{q})$ Equation 3.3

Gap-masked $Z(\mathbf{r})$ (forward and backward direction masking).

Synthetic data

$Q_{\mathrm{DW}}^{(i)}(\mathbf{r}) = h[\Delta(\mathbf{r})]$ Equation 4.3

Ideal inhomogeneous Q_{DW}.

$\mathcal{R}_s(\mathbf{r})$ Equation 4.9

Synthetic $\mathcal{R}(\mathbf{r})$.

A_s Section 4.4

	Amplitude of synthetic *d*-form factor modulations in $\mathcal{R}_s(\mathbf{r})$.	
ϕ_s	Phase of synthetic *d*-form factor modulations in $\mathcal{R}_s(\mathbf{r})$.	Section 4.4
A_{lat}	Amplitude of atomic corrugations in $\mathcal{R}_s(\mathbf{r})$.	Equation 4.9
f_{lat}	Spatial structure of atomic corrugations in $\mathcal{R}_s(\mathbf{r})$.	Equation 4.9
A_{noise}	Amplitude of white noise in $\mathcal{R}_s(\mathbf{r})$.	Equation 4.9
$\eta(\mathbf{r})$	Normally distributed white noise with unit standard deviation.	Equation 4.9
α	Atomic corrugation asymmetry parameter.	Equation 4.10

$\mathbf{Q}^{(\mathbf{i})}_{\mathrm{DW}}$ hypotheses

$h(\Delta)$	Hypothesis function to describe the doping-dependent charge modulation wavevector.	Equation 4.3
m_l	Slope of the linear hypothesis.	Figure 4.1

b_l	Figure 4.1

Intercept of the linear hypothesis.

m_p	Figure 4.1

Slope for the incommensurate portion of the piecewise hypothesis.

b_p	Figure 4.1

Intercept for the incommensurate portion of the piecewise hypothesis.

Δ_c	Figure 4.1

Within the commensurate to incommensurate transition model of Q_{DW}, the value of Δ at the transition.

δ_c	Section 4.7

(Notation for Section 4.7 only.) Within the commensurate to incommensurate transition model of Q_{DW}, the value of Δ at the transition.

$g_l(m_l, b_l)$	Section 4.7.1

Linear hypothesis probability distribution function.

$g_p(m_p, b_p)$	Section 4.7.1

Piecewise hypothesis probability distribution function.

Acknowledgments

The data analyzed in this thesis was collected by Michael Boyer (UD25K, UD32K, OPT35K) and Yi Yin (OD15K), on samples grown by Takeshi Kondo, Tsunehiro Takeuchi and Hiroshi Ikuta. Early analysis by Yang He inspired this deep look at the *d*-form factor density wave, and Chapter 3 is largely a result of discussions with Mohammad Hamidian and Jenny Hoffman. The idea to apply spatial masking grew out of discussions with Eric Hudson, and Debanjan Chowdhury contributed insight in interpreting the findings and collaborated with me writing Chapter 3. Thank you to Subir Sachdev for insightful discussions as we analyzed the data. Thank you also to Andrej Mesaros, Patrick A. Lee, Brad Ramshaw, Simon Verret and Milan Allan for providing feedback.

The project to apply machine learning to analyze the data was inspired by the work of Yi Zhang, Andrej Mesaros and Eun-Ah Kim, who have generously discussed ideas and shared insights with us. The work presented in Chapter 4 has been a close collaboration with Kaylie Hausknecht. Eric Hudson wrote the main density wave simulation algorithm. Jenny Hoffman proposed the project, and has continued to provide invaluable feedback throughout. Finally, in the last several months, this work had benefitted greatly from discussions with Trevor Rhone. Neural network computations were run on the Odyssey cluster supported by the FAS Division of Science, Research Computing Group at Harvard University.

I joined Jenny Hoffman's lab in 2014, and my work over the last five years has been shaped by her insightful and practical feedback. Jenny has pushed me to be more ambitious than I had ever intended by repeatedly sending opportunities my way and by demonstrating confidence in me since the day I started working with her. Her confidence in the abilities of her students and her prioritization of mentorship within the group are truly exceptional. I am grateful for her candid scientific and professional guidance and also for her sincere and unrelenting kindness.

My experience at Harvard has been shaped by working on the MBE-STM experiment in the Hoffman lab. Dennis Huang patiently trained me in 2014 and was a role model in a number of ways. Harry Pirie worked alongside me for years to troubleshoot a series of instrumentation and MBE growth challenges. Christian Matt joined the experiment with enthusiasm in 2017 and has worked tirelessly to make progress. I am grateful to both Christian and Harry for taking over responsibilities for me in the last year. I have had the opportunity to work with two dedicated undergraduate students Trevor Chistolini and Camille Bean. In addition, I have been lucky to work with Can-Li Song, Jason Hoffman, Ruizhe Kang, and Kevin Hauser.

It has been a pleasure (and really fun) to work with Kaylie Hausknecht on a project where we have both been learning constantly along the way. I was also fortunate to work with Mohammad Hamidian on both the work presented herein, as well as several other projects, and I have also benefitted from his advice along the way.

Over the last five years, I have enjoyed working with and getting to know many colleagues in the Hoffman lab: Yang He, Maoz Ovadia, Jianfeng Ge, Alyson Spitzig, Richard Liu, Eddie Ji, Larissa Little, Nathan Drucker, Ben November, Zhihuai Zhu, Pengcheng Chen, Oliver Yam, and Jessie Zhang. It has also been gratifying in the last year to be a part of a joint effort between the Hoffman and Mundy labs to expand MBE sample preparation facilities. I have enjoyed discussions and collaboration the Mundy group, including Julia Mundy, Charles Brooks, and Qi Song.

I would like to thank Dionne Clarke and Barbara Drauschke for the countless ways in which I have relied on them over the years. Thank you to Jacob Barandes and Lisa Cacciabaudo for making the PhD program as easy to navigate as possible.

I have also continued to benefit from the guidance of Leo Nikkinen, Zaven Altounian and David Ronis while I was at McGill (some of which, I certainly did not fully appreciate at the time), and I continue to benefit from the mentorship and support of my former advisor Dominic Ryan.

None of this work would have been possible without the support of family and friends. I am fortunate to have met a wonderful group of people through Rocket Team. Charles Xiong and Lucie Baumont have been like family to me here on the East Coast. Finally, I am grateful to my parents Randy and Lauren, and my sister Olivia for sharing in this journey with me and offering unwavering support.

1

Introduction

Condensed matter systems with strong electron-electron correlations exhibit a wide range of phenomena, including exotic quantum states such as spin liquids and the cuprate pseudogap phase, as well as technologically desirable properties such as colossal magneto resistance, metal-insulator transitions, and high temperature superconductivity [2, 3, 4, 5]. Strong interactions give rise to many-body behavior which cannot be described in the single-particle picture where electrons are filled into non-interacting (Fermi gas) or weakly-interacting (Fermi liquid) states. Strongly correlated materials have rich phase diagrams where a wide range of properties can be accessed by tuning the electronic structure through doping, interface engineering in heterostructures, pressure, etc...[5, 6, 7, 4, 3, 8]

Experimentally probing these complex phase diagrams can be challenging because (1) strong correlations produce a tendency for inhomogeneity [7], complicating the interpretation of bulk-averaged measurements, and (2) strong electron-electron repulsion favors localization of the electronic states in real-space (e.g. a Mott insulator [9]), whereas typical metals have itinerant states with sharp low energy quasiparticle excitations in crystal momentum k-space. To understand conducting strongly correlated materials can therefore require a combination of real and momentum space probes. Scanning tunneling microscopy (STM) [10], capable of imaging electronic structure with angstrom-scale spatial resolution and whose data can be analyzed in both real and momentum space, enables disentangling the complex spatial structure and characterizing both localized and itinerant electronic properties.

1.1 Cuprate superconductors

In cuprates, high temperature superconductivity emerges from doping (with electrons or holes) between an antiferromagnetic (AFM) insulating phase, where strong electron repulsion localizes the Cu $d_{x^2-y^2}$ electrons and a metal where the kinetic energy favors delocalization [2, 7]. We restrict the remainder of this discussion, and this thesis, to the hole-doped cuprates. In addition to unconventional superconductivity, the strong electronic interactions give rise to a variety of exotic electronic behaviors. It has long been proposed that a quantum critical point (QCP), separating the so-called pseudogap (PG) phase from the metallic phase resembling a Fermi liquid, may lie at the center of the phase diagram [11] (Figure 1.1). While evidence supporting the quantum critical scenario has continued to emerge [12, 13, 14, 15], the existence and nature of the putative QCP remain open questions.

While a number of broken symmetries have now been observed, none to date provide a satisfactory description of the PG. Recently, inversion [17] and rotational [14, 13] symmetries were reported to be broken along a line of phase transitions [15, 17, 14, 13] bounding the PG state, but alone cannot explain the spectral gap. Alternatively, the onsets of finite **Q** broken symmetries, spin and charge order, do not match the PG onset [6, 18, and references therein]. There are proposals that pair density wave order [19, 20, 21], a modulation of the Cooper pair density associated with finite **Q** pairing, may underlie the PG, and recent experiments [22, 23] support the existence of Cooper pair modulations. However, there are also proposals that fluctuating rather than static order may give rise to the PG [24, and references therein] and alternatively, that the transition may be beyond the Ginzburg-Landau paradigm [25, and references therein].

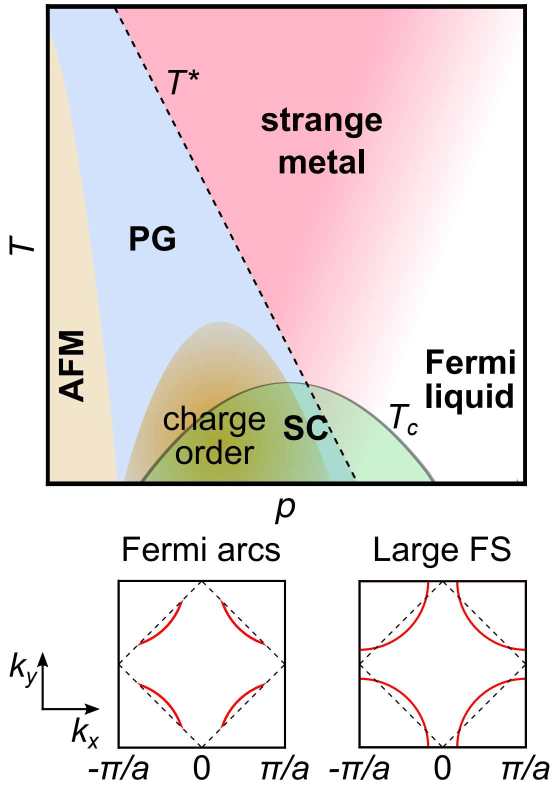

Figure 1.1: Schematic phase diagram and Fermi surface (FS) structure for cuprate high temperature superconductors (largely based on the phase diagrams presented in References [6, 16]).

1.2 Materials and crystal structure

In 1986, Bednorz and Müller discovered superconductivity up to 30 K in cuprate $Ba_xLa_{5-x}Cu_5O_{5(3-y)}$ [1]. Subsequently, Y-, Bi-, Tl- and Hg-based compounds were discovered, raising T_c above the temperature of liquid nitrogen. The record for the highest T_c under ambient pressure is 135 K in $HgBa_2Ca_2Cu_3O_{1+x}$ [26, 27]. Adjusting the concentration of holes p (typically via chemical substitution and oxygen content) tunes T_c such that the superconducting phase has a dome shape in the temperature-doping phase diagram (Figure 1.1). Compounds are referred to as overdoped (OD) or underdoped (UD) in comparison to the doping with optimal T_c and are typically named by the OD, UD or OPT designation with the value of T_c, *e.g.* UD25K. The high T_c and the symmetry of the superconducting order parameter (d- rather than s-wave) [28] were early indications of the unconventional pairing mechanism.

All cuprate superconductor families share a layered structural motif with the low energy electronic structure arising from the Cu $d_{x^2-y^2}$ and O p orbitals in CuO_2 layers that have a two-dimensional square lattice structure. The number of CuO_2 planes per unit cell and the composition of the spacer layers vary between families.

The Bi-based cuprates have been studied the most heavily by surface sensitive scanning tunneling microscopy (STM) and angle resolved photoemission spectroscopy (ARPES) because they have a non-polar cleavage plane between BiO layers. This thesis focuses on the family $(Bi,Pb)_2(Sr,La)_2CuO_{6+\delta}$ (Bi2201, Figure 1.2), a Bi-based

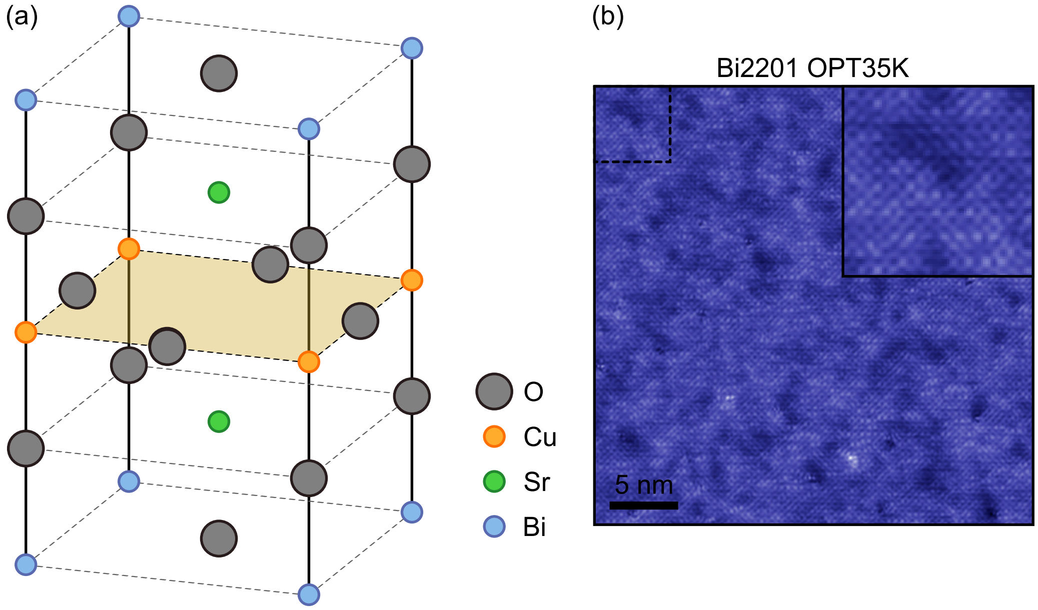

Figure 1.2: **(a)** Bi2201 crystal structure. **(b)** Atomically resolved STM topography of a Bi2201 sample near optimal doping with 35 K T_c. The inset shows an enlargement of the region enclosed by the dashed box.

cuprate with a single CuO_2 plane in the unit cell. While Bi2201 has a lower optimal T_c ($\approx$ 35 K [29]) than the closely related double layer compound $Bi_2Sr_2CaCu2O_{8+x}$ (Bi2212) (92 K [30, 31]), analysis of its electronic structure is simpler due to the absence of bilayer splitting. Furthermore, Pb-doping suppresses the structural distortion called supermodulation that affects many cuprate families [32].

1.3 Pseudogap

The PG phase is a phenomenologically defined region of the phase diagram where the electronic density of states has an unexplained spectral gap (the pseudogap) at the Fermi level. This state is also the normal state from which superconductivity emerges in the underdoped and optimally doped regimes. The PG was first discovered from nuclear magnetic resonance (NMR) measurements [33, 34] which show a suppression of both the Knight shift (proportional to the real part of the paramagnetic susceptibility) and the spin lattice relaxation rate (probes the imaginary part of the susceptibility). Additional spectroscopic probes, including ARPES, tunneling and Raman spectroscopy have confirmed the gap with onset at a temperature T^* well above the superconducting T_c[†]. Both T^* and the energy scale of the PG are maximum at low doping and decrease with p until near optimal doping they become difficult to distinguish from T_c and the energy scale of coherent superconducting pairing, respectively [36, 16, 37]. ARPES has shown that similarly to the SC gap, the PG is largest along the $(k_x, 0)$ and $(0, k_y)$ directions, giving rise to an unexplained Fermi surface structure seemingly consisting of open arcs [38, 39, 40, 41, 42].

1.4 FERMI SURFACE STRUCTURE

Understanding the Fermi arcs remains one of the largest puzzles in the cuprate phase diagram. In the OD Fermi liquid region, ARPES and quantum oscillations measure a hole-like Fermi surface centered at the corner of the Brillouin zone with a conventional area of $1 + p$ [43, 44] consistent with Luttinger's theorem [45, 46]. However, moving to lower doping, the antinodal (along the direction of the superconducting

---

[†] A number of reviews cover the pseudogap phenomenology, for example, References [35, 2, 36, 6, 16].

gap maximum) segments of the Fermi surface become gapped, leaving only the Fermi arcs [38, 39, 40, 41, 42]. Cooling below T_c, Bogoliubov quasiparticles in the superconducting state do not disperse all the way to the edge of the Brillouin zone, but terminate at the ends of the normal state arcs [47, 48, 49].

In Bi2201, a spectral gap persists out to the overdoped edge of the superconducting dome [50, 51], but at low temperature, sharp antinodal quasiparticles in the superconducting state vanish near optimal doping [47]. Because both features are typically associated with the pseudogap, for clarity, we refrain from using the language of the pseudogap in the remainder of this work, and instead refer directly to the feature (Fermi arc or spectral gap) of interest.

1.5 Density wave

Charge density modulations have been observed ubiquitously across underdoped cuprates [18]. Charge order was first observed by neutron scattering as part of stripe order where doped holes segregate at domain boundaries of antiferromagnetic order (forming stripes) in Nd-substituted LSCO ($La_{2-x-y}Nd_ySr_xCu_2O_{4+\delta}$) [52]. In 2002, Hoffman *et al.* [53] imaged four-unit-cell periodic modulations of the local density of states surrounding vortex cores in Bi2212 using STM, and in 2003, Howald *et al.* [54] imaged modulations under zero field. Subsequently, similar modulations have been characterized not only by STM, but also resonant x-ray scattering, x-ray diffraction and neutron scattering[*] in La-, Y-, Bi-, Hg- and Na-based families [18]. Recent advances in STM enabling sublattice-resolved analysis of the charge modulation structure have indicated that the intra-unit cell structure in Bi2212 and NaCCOC is primarily d-form factor (dFF), meaning that the modulations are localized on oxygen sites, with the “waves” on the two oxygen sublattices out of phase [55, 56], as will be discussed in Section 2.3. Resonant x-ray scattering measurements reach the same conclusion for YBCO [57, 58], were inconclusive for Bi2201 [57], and indicate stripe order in LBCO primarily has an s'-form factor [58].

There are many proposals for the origin(s) of the observed charge modulations. Microscopic models can be broadly grouped based on whether they treat the order as arising from strong real space interactions [59, 7, and references therein], or a momentum-space instability [60, 61, 62, 63]. The vast majority of the reported modulation wavevectors are incommensurate [18], and have been interpreted as tied to the structure of the Fermi surface [64, 65]. However, the short correlation length makes measurement of the underlying wavevector challenging [66], and with the exception of Wise *et al.* [64], the measurements have not accounted for spatial inhomogeneity. Recent work on Bi2212 has found that the wavevector over a wide doping range is consistent with a commensurate four-unit-cell modulation [66], but it is unclear how to unify the picture across cuprate families.

---

[*] Neutron scattering does not directly measure the charge order, but measures the coupled ionic displacements.

Ubiquity of charge ordering in the underdoped cuprates has spawned speculation of an underlying relationship to the pseudogap, and has been used to motivate intense research efforts. STM has shown that the intensity of dFF modulations is maximized near the spectral gap energy, Δ, and the ordering wavevector is similar to the wave vector connecting the Fermi arc tips [64, 65, 67]. Furthermore, the onset of charge order matches the transition from large Fermi surface to Fermi arcs in Bi2212 [48]. But the latter relationship is not universal across cuprate families [68, 69]. Concrete evidence, ideally the observation of coordinated responses in Fermiology and charge order under the effect of a controlled tuning parameter, is necessary to substantiate a link between the two phenomena and to uncover the nature of such a link.

1.6 ELECTRONIC INHOMOGENEITY

Doping via chemical substitution and oxygen content leads to severe inhomogeneity in cuprates [70, 64, 71, 72, 73, 74, 75]. Scanning tunneling microscopy has studied this inhomogeneity intensely, and has shown that the electronic structure varies on nanometer length scales and can be correlated to the local concentration of individual dopants [73, 70, 72]. The spatial evolution of features which are associated with the doping evolution across the phase diagram—the spectral gap Δ, the Van Hove singularity [74], the charge order wavevector [64], and the Fermi surface structure [64]—suggests that the doping is determined locally.

1.7 Goal of this work and outline

In this work, we set out to elucidate the nature of the ground state reorganization that occurs near optimal doping. We use dFF charge modulations in Bi2201 imaged via scanning tunneling spectroscopy as a doping-dependent fingerprint of the underlying electronic correlations. Furthermore, the nanoscale inhomogeneity means that a single STM image of the disordered modulation structure encodes the doping-dependence over a region of the phase diagram.

We proceed as follows. In Chapter 2, we describe interpretation of STM data and relevant analysis techniques. In Chapter 3, we disentangle the spatial inhomogeneity to uncover a commensurate to incommensurate transition in the dFF charge order. Finally, in Chapter 4, we address the challenges associated with analyzing the complex spatial structure by applying convolutional neural networks to distinguish between different ordering scenarios.

2

Experimental methods

2.1 WHAT STM MEASURES

Extensive literature about the operational principles of a scanning tunneling microscope (STM), and the theory of STM can be found elsewhere.[*] Here, we summarize the principles necessary to understand the work presented in this dissertation.

---

[*] See, for example, the text by Chen [76].

Scanning tunneling microscopy enables atomic-scale imaging of the local density of states (DOS) at the surface of a material. An atomically-sharp metallic tip, which can be displaced in all three spatial dimensions with pico-scale precision, is brought close (< 1 nm) to the sample surface, enabling electrons to tunnel across the junction. Applying a bias voltage V_{bias} across the junction (here applied to the sample) produces a tunneling current I. At low temperature, when V_{bias} is positive, on average electrons tunnel from the tip to the sample, and I is dominated by an integration over contributions of sample states from the Fermi energy E_F to $E_F + eV_{\text{bias}}$, where e is the electron charge. Likewise, for negative V_{bias}, with contributions primarily coming from the filled sample states. The differential conductance $g(\mathbf{r}, E) = dI/dV(\mathbf{r}, E)$, where $\mathbf{r}$ is the two-dimensional position of the tip on the sample surface, is proportional to the local DOS. Under the simplifying conditions of low temperature providing a sharp Fermi distribution cutoff, and a tip with DOS constant in energy, the tunneling current and differential conductance are approximated by:

$$\begin{aligned} I(\mathbf{r}, V) &\approx \frac{4\pi e}{\hbar} e^{-z\sqrt{\frac{8m\varphi}{\hbar^2}}} N_t \int_0^{eV} N(\mathbf{r}, \varepsilon)\, d\varepsilon && (2.1) \\ g(\mathbf{r}, E) &\propto N(\mathbf{r}, E), && (2.2) \end{aligned}$$

where we have set $E_F = 0$, z is the tip-sample separation, m is the electron mass, φ is the tunneling barrier height, N_t is the tip DOS and N is the sample DOS. Finite temperature introduces a spread in energy, and while the STM tip is prepared to have constant DOS, the tip will generically induce an energy dependent background.

Scanning the tip across the surface enables spatially mapping the local DOS, but $g(\mathbf{r}, E)$ contains a spatially varying prefactor. Typically, z is adjusted by a feedback loop to reach a specified current setpoint I_{s}. If we then keep z fixed but change V from the bias V_{s} used to set up the junction, as we typically do for spectroscopy,

$$
\begin{aligned}
I(\mathbf{r}, V) &\approx \frac{I_{\mathrm{s}}}{\int_0^{eV_{\mathrm{s}}} N(\mathbf{r}, \varepsilon)\, d\varepsilon} \int_0^{eV} N(\mathbf{r}, \varepsilon)\, d\varepsilon & (2.3) \\
g(\mathbf{r}, E) &\approx \frac{eI_{\mathrm{s}}}{\int_0^{eV_{\mathrm{s}}} N(\mathbf{r}, \varepsilon)\, d\varepsilon} N(\mathbf{r}, E), & (2.4)
\end{aligned}
$$

The spatially varying prefactor depends on the setup conditions. The ratio map $Z(\mathbf{r}, E) = g(\mathbf{r}, E)/g(\mathbf{r}, -E)$ corrects for this measurement artifact called the setup effect [77]. Furthermore, due to particle-hole mixing in the superconducting state, the Z ratio map, which can be expressed in terms of the BCS factors $Z(\mathbf{r}, E) \approx |u(\mathbf{r}, E)|^2/|v(\mathbf{r}, E)|^2$ enhances the signal from Bogoliubov quasiparticle interference used to infer the normal state Fermi surface structure [78, 79].

When examining charge order in Chapter 3, we will use one final quantity, $\mathcal{R}(\mathbf{r}, E)$:

$$
\mathcal{R}(\mathbf{r}, E) = \frac{\sum_{0<E_i<E} g(\mathbf{r}, E_i)}{\sum_{-E<E_i<0} g(\mathbf{r}, E_i)} \qquad (2.5)
$$

The sums enhance the signal originating from the non-dispersive density wave, which spreads over a wide range in energy range. The ratio, as with Z, suppresses the setup effect. In other works, the notation $\mathcal{R}(\mathbf{r})$ is often used to denote a similar quantity calculated from I rather than g as $\mathcal{R}(E) = I(E)/I(-E)$. Under typical assumptions, the two versions of $\mathcal{R}$ are equivalent to the degree that the sums in Equation 2.5 approximate the integrals in $I(E) = \int_{E_F}^{E} g(\varepsilon)\, d\varepsilon$.

2.2 Data examined in this work

Differential conductance maps were collected in scanning tunneling microscopes at 6 K with the following tip-sample junction setup conditions: 100 mV and 100 pA for OD15K, -100 mV and 400 pA for OPT35K, -200 mV and 400 pA for UD32K, and −150 mV and 400 pA for UD25K. Measurements used a lock-in technique with 2 mV, 10 mV, 5 mV, and 5 mV bias modulations, respectively. Data were corrected for artificial distortions due to instrument drift to register the simultaneously recorded topography to a perfect lattice [77].

In the following sections, we examine in detail data processing to access the structure of charge order present in the data. The glossary at the beginning of this dissertation summarizes all of the relevant quantities characterizing and derived from STM measurements.

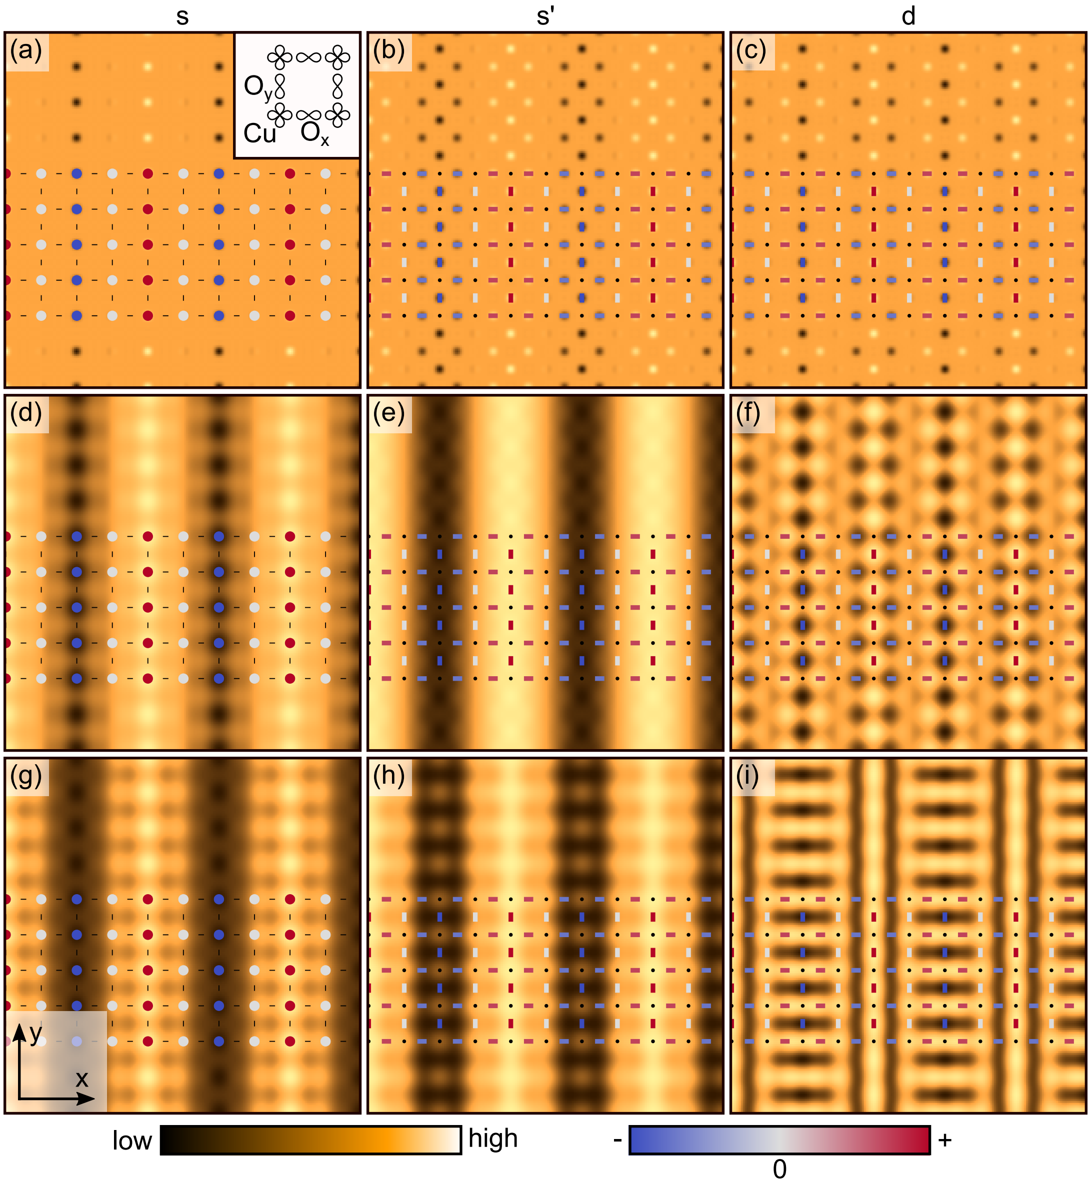

Figure 2.1: Charge modulation intra-unit cell structure. (a–c) Simulated s-, s'- and d-form factor modulations, using momentum cutoff parameter $w = 2.0$ r.l.u. Schematic representations of the modulations on the lattice are overlaid, with Cu sites represented by circles, and O sites represented by bars. The inset in (a) is a schematic of the CuO_2 unit cell. **(d–f)** Simulated s-, s'- and d-form factor modulations, using momentum cutoff parameter $w = 0.5$, which is closer to the bandwidth of our STM data. **(g–i)** Modulations shown in (d–f) with atomic corrugations added in.

2.3 Form factor components

The unit cell of the CuO_2 plane contains one Cu atom at $(0,0)$ and two O atoms, O_x at $(a/2,0)$ and O_y at $(0,a/2)$ (a is the lattice constant), meaning that charge modulations can display non-trivial intra-unit cell structure, depending on the nature of the local electronic correlations. The intra-unit cell structure can be described in terms of angular momentum form-factors: an s form factor (sFF) modulation is localized on the Cu sites, whereas s' and d form factors are localized on the O sites: O_x and O_y modulations are in phase for s' and out of phase for d (shown in Figure 2.1) [55, 80, 81, 61, 57]. Concretely, charge modulations can be expressed in terms of form factor components:

$$
\begin{aligned}
\rho_s(\mathbf{r}) &= \begin{cases} \mathcal{A}_s \cos(\mathbf{Q}\cdot\mathbf{r}+\varphi_s) & \text{for } \mathbf{r}=\mathbf{r}_{\mathbf{n},\mathrm{Cu}} \\ 0 & \text{for } \mathbf{r}=\mathbf{r}_{\mathbf{n},\mathrm{O}_x} \\ 0 & \text{for } \mathbf{r}=\mathbf{r}_{\mathbf{n},\mathrm{O}_y} \end{cases} \\
\rho_{s'}(\mathbf{r}) &= \begin{cases} 0 & \text{for } \mathbf{r}=\mathbf{r}_{\mathbf{n},\mathrm{Cu}} \\ \mathcal{A}_{s'} \cos(\mathbf{Q}\cdot\mathbf{r}+\varphi_{s'}) & \text{for } \mathbf{r}=\mathbf{r}_{\mathbf{n},\mathrm{O}_x} \\ \mathcal{A}_{s'} \cos(\mathbf{Q}\cdot\mathbf{r}+\varphi_{s'}) & \text{for } \mathbf{r}=\mathbf{r}_{\mathbf{n},\mathrm{O}_y} \end{cases} \\
\rho_d(\mathbf{r}) &= \begin{cases} 0 & \text{for } \mathbf{r}=\mathbf{r}_{\mathbf{n},\mathrm{C}u} \\ \mathcal{A}_d \cos(\mathbf{Q}\cdot\mathbf{r}+\varphi_d) & \text{for } \mathbf{r}=\mathbf{r}_{\mathbf{n},\mathrm{O}_x} \\ -\mathcal{A}_d \cos(\mathbf{Q}\cdot\mathbf{r}+\varphi_d) & \text{for } \mathbf{r}=\mathbf{r}_{\mathbf{n},\mathrm{O}_y} \end{cases}
\end{aligned}
\tag{2.6}
$$

where $\mathcal{A}_{s,s',d}$ and $\varphi_{s,s',d}$ are the amplitude and phase of the corresponding modulations, and $\mathbf{r}_{\mathbf{n},\mathrm{Cu}}$, $\mathbf{r}_{\mathbf{n},\mathrm{O}_x}$ and $\mathbf{r}_{\mathbf{n},\mathrm{O}_y}$ are the positions of the corresponding lattice sites in the crystal unit cell indexed by $\mathbf{n}$. When $\mathbf{Q}=0$, the point group symmetry of each of these modulations is well-defined [55].

The three form factor components have distinct momentum space structures (Figure 2.2) which can be understood as follows. The expressions in equation 2.6 can each be rewritten as a product $\rho_l(\mathbf{r}) = a_l(\mathbf{r})\mathcal{A}_l \cos(\mathbf{Q}\cdot\mathbf{r}+\varphi_l)$, where l indexes the form factor, of a modulation $\mathcal{A}_l \cos(\mathbf{Q} \cdot \mathbf{r} + \varphi_l)$ with a function $a_l(\mathbf{r})$ describing the intra-unit cell structure, a sum over delta functions located at atomic positions. The Fourier transforms are therefore convolutions of the Fourier transformed modulation (peaks at $\pm\mathbf{Q}$) with the Fourier transform of the sublattice components. The Fourier transform of the Cu sublattice produces a sum over delta functions located at each reciprocal lattice vector $\mathbf{G}_{ij} = i\mathbf{G}_{\mathbf{x}} + j\mathbf{G}_{\mathbf{y}}$, such that the Fourier transform $\tilde{\rho}_s$ has a set of satellite peaks displaced by $\pm\mathbf{Q}$ from each $\mathbf{G}_{ij}$. The sums over O_x and O_y sublattices in the s'FF and dFF components (the O_y sublattice has a factor of -1 for d) each contain contributions from only half of the reciprocal lattice vectors and with alternating signs, leading to the patterns shown in Figure 2.2 b and c. The dFF component, which will be the focus of this work, has alternating signs between the peaks located around the x and y Bragg peaks at $\mathbf{G}_x$ and $\mathbf{G}_y$.

The spatial resolution of the STM and the spatial density of measurement points in the data imposes a bandwidth on the wavevectors whose oscillations can be resolved in the real-space image (*i.e.* a $\mathbf{q}$ cutoff or bandwidth). The cutoff from the pixel density of the UD32K data (the data set with the lowest spatial resolution, corresponding to the smallest $\mathbf{q}$ cutoff) is shown in Figure 2.2. A smooth momentum cutoff has been applied in the real-space simulations in Figure 2.1 as a Gaussian envelope with standard deviation w on the Fourier components of the sublattice structure. The top row (a–c) has a large bandwidth $w = 2$ r.l.u. (1 r.l.u. $= 2\pi/a$) to make the sublattice structure clear, but the middle row (d–f) with $w = 0.5$ r.l.u. is a more realistic expectation

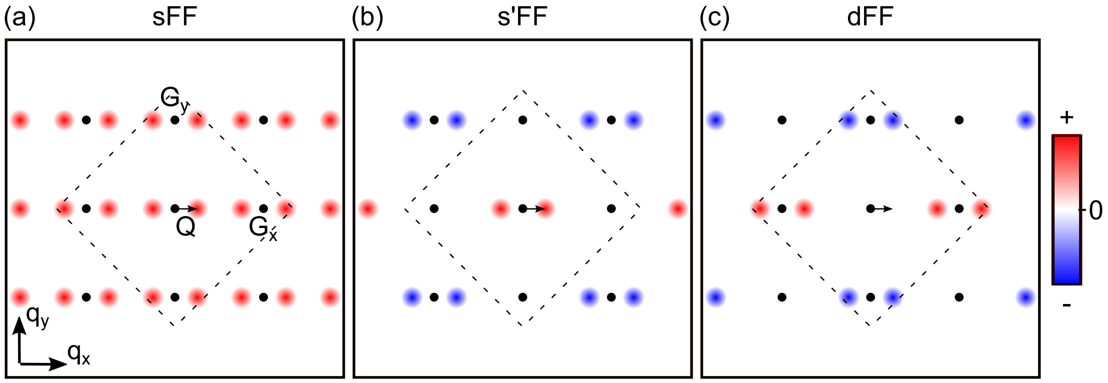

Figure 2.2: Form factor components in Fourier space. Schematics of the Fourier transforms of s-, s'- and d-form factor modulations. $\mathbf{Q}$ labels the modulation wavevector, $\mathbf{G}_x$ and $\mathbf{G}_y$ label reciprocal lattice vectors (the locations of Bragg peaks), and the dashed line indicates the smallest $\mathbf{q}$ bandwidth of the data sets examined in this work (UD32K).

of the structures that we search for in the STM data.

We follow Fujita et al. [55] in using intra-unit cell masking to extract specific form factor components of $\mathcal{R}(\mathbf{r})$ (Figure 2.3). The general idea is to selectively isolate a specific form factor by masking individual sublattices. First, we generate masks for each sublattice:

$$\begin{aligned} M_{\mathrm{Cu}}(\mathbf{r_{ij}}) &= \begin{cases} 1.0 & \text{for } \mathbf{r}_{ij} = \mathbf{r}^{\mathrm{px}}_{\mathrm{Cu}} \\ 0.0 & \text{otherwise} \end{cases} \\ M_{\mathrm{O}_x}(\mathbf{r_{ij}}) &= \begin{cases} 1.0 & \text{for } \mathbf{r}_{ij} = \mathbf{r}^{\mathrm{px}}_{\mathrm{O}_x} \\ 0.0 & \text{otherwise} \end{cases} \\ M_{\mathrm{O}_y}(\mathbf{r_{ij}}) &= \begin{cases} 1.0 & \text{for } \mathbf{r}_{ij} = \mathbf{r}^{\mathrm{px}}_{\mathrm{O}_y} \\ 0.0 & \text{otherwise} \end{cases} \end{aligned} \tag{2.7}$$

where $\mathbf{r_{ij}}$ is the position of pixel indexed by i and j, and the px superscript on the atom positions indicates the nearest pixel position. For simplicity of notation, the unit cell index $\mathbf{n}$ has been dropped. The form factor components of $\mathcal{R}(\mathbf{r})$ are then isolated as:

$$\begin{aligned} S(\mathbf{r_{ij}}) &= M_{\mathrm{Cu}}(\mathbf{r_{ij}})\mathcal{R}(\mathbf{r_{ij}}) \\ S'(\mathbf{r_{ij}}) &= \left[M_{\mathrm{O}_x}(\mathbf{r_{ij}}) + M_{\mathrm{O}_y}(\mathbf{r_{ij}})\right]\mathcal{R}(\mathbf{r_{ij}}) \\ D(\mathbf{r_{ij}}) &= \left[M_{\mathrm{O}_x}(\mathbf{r_{ij}}) - M_{\mathrm{O}_y}(\mathbf{r_{ij}})\right]\mathcal{R}(\mathbf{r_{ij}}) \end{aligned} \tag{2.8}$$

It is crucial that the data is first corrected for drift via the Lawler-Fujita algorithm [77], such that the sublattice positions can be determined accurately.

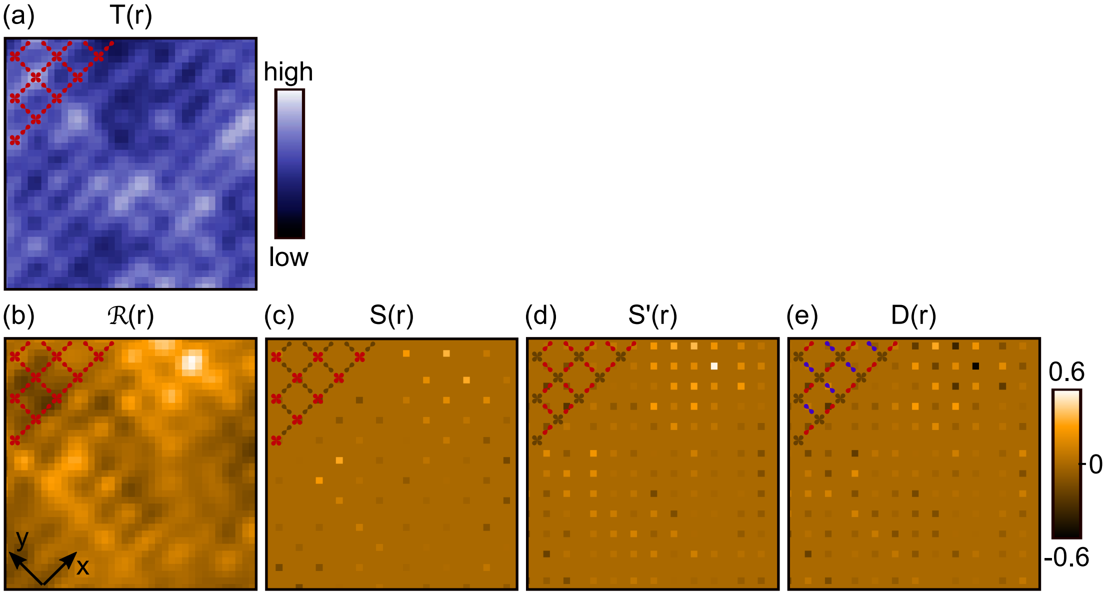

Figure 2.3: Form factor decomposition in UD32K. (a) Topography from UD32K. **(b)** $\mathcal{R}(\mathbf{r})$ from UD32K. **(c–e)** Form factor components of (b). All panels show the same field of view, and schematics of the lattice are overlaid to illustrate the form factor masks. Red and blue indicate +1 and −1 mask values.

2.4 EXAMINING THE DENSITY WAVE STRUCTURE

In the last section, we focused on the form factor $a_l(\mathbf{r})$ of a modulation $a_l(\mathbf{r})\mathcal{A}_l\cos(\mathbf{Q}\cdot\mathbf{r}+\varphi_l)$. Now we focus on the cosine modulation, and consider how the structure changes when generalizing to include disorder. This section is largely based on work by Mesaros *et al.* [66].

In general, a disordered wave can be expressed as

$$\psi(\mathbf{r}) = A(\mathbf{r})\cos[\phi(\mathbf{r})], \tag{2.9}$$

where a perfect wave with wave vector $\mathbf{Q}$ has linear phase $\phi(\mathbf{r}) = \mathbf{Q}\cdot\mathbf{r}+\varphi$ and constant amplitude $A(\mathbf{r}) = A$. However, the charge modulations in the cuprates are severely disordered, with correlation lengths of approximately $11a$ (just over three wavelengths) in $YBa_2Cu_3O_{6.6}$ [82], $6a$ (less than two wavelengths) in $Bi_2Sr_2CaCu_2O_{8+x}$ [83], and $4a$ (roughly one wavelength) in Bi2201 [65]. The disorder represented by the spatial variations in $A(\mathbf{r})$ and $\phi(\mathbf{r})$ are therefore simultaneously important – encoding the sources of the disorder – and difficult to extract. In this section, we address how to determine the spatially varying amplitude and phase from an image of the modulation.

Let us start by motivating how information is encoded in the modulation. First, note that the local wavevector of the modulation is the gradient of the phase $\nabla\phi(\mathbf{r})$.

Then consider two causes of charge density wave (CDW) inhomogeneity:

1. In typical charge density wave systems where an incommensurate instability competes with an energy term favoring commensuration with the lattice, it is well known that it becomes favorable to introduce discommensurations, soliton defects in the phase, in order to increase registry to the lattice [84]. The wave has multiple degenerate configurations corresponding to shifting by an integer number of crystal unit cells. The phase increases linearly over regular intervals interrupted by abrupt phase slips between degenerate configurations, and the phase slips add up such that the average wave vector deviates from the wavevector within each domain.
2. A similar structure can arise in commensurate CDWs. Pinning by defects can cause domains of the density wave to be separated by discommensurations with random phase slips at the boundaries such that the average wavevector matches the local wavevector within the domains.

In each of these cases, we can write the modulation as $\psi(\mathbf{r}) = A\cos[\mathbf{Q}\cdot\mathbf{r} + \varphi_{\mathrm{D}}(\mathbf{r})]$, where φ_{D} contains the domain structure of the wave, *i.e.* it has a constant value within each domain reflecting the phase of the wave with respect to the reference wave $\cos(\mathbf{Q}\cdot\mathbf{r})$, but changes across domain walls.

We access the local phase and amplitude variations in the demodulated wave, calculated in Fourier space:

$$\tilde{\Psi}_{\mathbf{Q}}(\mathbf{q}) = \exp\left[-\frac{q^2}{2\Lambda^2}\right]\tilde{\psi}\left(\mathbf{q}+\mathbf{Q}\right), \tag{2.10}$$

where the tilde indicates the Fourier transform, and $\tilde{\Psi}_{\mathbf{Q}}(\mathbf{q}) = 0$ outside of a domain that isolates the modulation's peak in q-space. In Chapters 3 and 4, we will use $\tilde{\psi}(\mathbf{q}) = \tilde{D}(\mathbf{q})$, the dFF component of the $\mathcal{R}$ map, and chose the domain of $\tilde{\Psi}_{\mathbf{Q}}(\mathbf{q})$ to isolate the dFF DW peak. Note that we use uppercase Ψ and Φ to differentiate the demodulated wave from the original wave ψ with spatial phase ϕ. In addition, Λ provides a smooth cutoff.[†] The inverse Fourier transform is a complex-valued function:

$$\Psi_{\mathbf{Q}}(\mathbf{r}) = A_{\mathbf{Q}}(\mathbf{r}) \exp[i\Phi_{\mathbf{Q}}(\mathbf{r})] \tag{2.11}$$

The phase $\Phi_{\mathbf{Q}}(\mathbf{r})$ is referenced to $\mathbf{Q} \cdot \mathbf{r}$ because of the shift by $\mathbf{Q}$ in Equation 2.10 and approximates $\phi(\mathbf{r}) - \mathbf{Q} \cdot \mathbf{r}$. Hence, for a perfect cosine wave, $\Phi_{\mathbf{Q}}(\mathbf{r})$ is constance, and for the two examples of domains separated by discommensurations, $\Phi_{\mathbf{Q}}(\mathbf{r})$ approximates φ_{D}, giving the domain structure. Alternatively, if the modulation has a different underlying wavevector, *i.e.* $\psi(\mathbf{r}) = A\cos[\mathbf{Q}_1 \cdot \mathbf{r} + \varphi_{\mathrm{D}}(\mathbf{r})]$, then the phase of $\Psi_{\mathbf{Q}}(\mathbf{r})$, where $\mathbf{Q} \neq \mathbf{Q}_1$, will change throughout the domain. With correct choice of $\mathbf{Q}$, the amplitude $A_{\mathbf{Q}}(\mathbf{r})$ approximates the spatially varying modulation amplitude $A(\mathbf{r})$.

Finally, we address the challenge of determining the correct wavevector $\mathbf{Q}$ describing the modulation. Mesaros et al. [66] pointed out that the standard technique of fitting the peak in Fourier space is sensitive to strong disorder, and instead proposed demodulation residue minimization as a well-defined alternative. The demodulation

---

[†] These hard and soft cutoffs isolate the dFF modulation of interest from the other signals contained in the STM data.

residue is defined by:

$$\begin{aligned} R_{\mathbf{Q}}[\psi] &= \sqrt{\left(R^{x}_{\mathbf{Q}}[\psi]\right)^2 + \left(R^{y}_{\mathbf{Q}}[\psi]\right)^2} \\ R^{\alpha}_{\mathbf{Q}} &\equiv \int \frac{d^2r}{L^2}\,\mathrm{Re}\left[\Psi^{*}_{\mathbf{Q}}(\mathbf{r})\,(-i\partial_\alpha)\,\Psi_{\mathbf{Q}}(\mathbf{r})\right] \end{aligned} \tag{2.12}$$

where $\alpha = x, y$ and L is the size of the data's field of view (square).

The derivative in Equation 2.12 acts on the phase, giving $|A_{\mathbf{Q}}|^2\,\partial_\alpha\Phi_{\mathbf{Q}}(\mathbf{r})$ inside the integral. If the amplitude is constant, then the residue corresponds to an integral over $\partial_\alpha\Phi_{\mathbf{Q}}(\mathbf{r})$, approximating the wavevector mismatch $\nabla\phi - \mathbf{Q}$, such that R is minimized with respect to $\mathbf{Q}$ when $\mathbf{Q}$ is the average wavevector.[‡] For real waves, the amplitude factor inside the integral does two things. First, in physical systems, the amplitude goes to zero at singularities where $\nabla\phi$ is undefined, keeping the residue well-defined. Second, in regions of low DW amplitude, the phase of the modulation is not well-defined by the data, and the amplitude factor avoids having these regions dominate the wavevector measurement.

[‡] This statement depends on being able to appropriately isolate the modulation of interest and an appropriate choice of Λ.

3

Density wave probes cuprate quantum phase transition

In cuprates, high-temperature superconductivity lies between an undoped antiferromagnetic (AFM) insulator and a metal at high hole doping (p). In proximity to the AFM insulator, the strong electronic correlations give rise to a complex phenomenology, including a large spectral gap Δ that opens above T_c, and a k-space structure lacking a conventional Fermi surface (FS) but described by open arcs [2, 6, 38, 39, 40, 41, 42]. Both gap and arcs are widely considered hallmarks of this underdoped region of the phase diagram, and have drawn significant attention aimed at uncovering their origin(s) [*]. However, at a doping near optimal superconductivity, the Fermi arcs undergo an abrupt transition to a 'large' pocket consistent with a conventional area proportional to $1 + p$ [43, 85, 68, 86]. A crucial challenge remains to identify the appropriate ground state(s) that underlie the theoretical framework on *both* sides of this transition.

On the overdoped side, long thought to be a Fermi liquid, several recent reports of anomalous behavior call into question the conventional interpretation [87, 43, 88, 89, 90, 91]. The observations of resistivity linear in temperature [88, 87] in Bi-based and La-based families challenge the expectations of standard Fermi liquid theory, and in MBE-grown $\mathrm{La_{2-x}Sr_xCuO_4}$ compounds, there are reports of mysterious symmetry breaking [90] and anomalous scaling of the superfluid density with critical temperature [89, 91], although the latter remains controversial [92, 93]. Furthermore, resonant inelastic x-ray experiments on Tl-, Y- [94], and La-based [95] compounds revealing spin fluctuations have been interpreted in terms of significant electron correlations, and there are now theoretical proposals [96, 97] for how overdoped compounds may

---

[*] The word pseudogap has been variously used to describe both spectral gap and Fermi arc phenomena; here we note the distinction between these phenomena.

retain certain characteristic features of Fermi-liquid-like behavior, while exhibiting fractionalization in the presence of strong stripe fluctuations. Additional experiments are necessary to understand the overdoped compounds, and in particular to clarify the extent to which the effects of strong correlations may persist through the FS transition.

In nearly all cuprate families, charge modulations have been reported in underdoped compounds, with detection terminating at [48] or before [69, 68] the doping where the FS transition occurs. In $(\mathrm{Bi,Pb})_2(\mathrm{Sr,La})_2\mathrm{CuO}_{6+\delta}$, Bi2201, however, charge modulations extend into the overdoped regime [47, 98]. These modulations, reflecting an ordering instability of the electronic system, therefore serve as a doping-dependent fingerprint of underlying electronic interactions, not just in the underdoped regime, but across the FS transition (Figure 3.1a). In the Bi2201 phase diagram, the FS transition occurs just below optimal doping [†], while the spectral gap persists in the presence of the large FS [47, 50, 51]. Peng *et al.* [98] recently investigated the charge modulation structure across the closing of the gap in the far-overdoped region, but it remains crucial to clearly define how the wavevector, Q_{DW}, evolves across the FS transition (Figure 3.1a). Existing measurements in this doping range [65, 99, 100], exhibit large scatter and doping coverage insufficient to clearly establish the Q_{DW} trend

---

[†] To the authors' knowledge, the FS transition in Bi2201, near $p = 0.14$, has been reported only by scanning tunneling microscopy (STM) [47], but Kondo *et al.*'s ARPES measurements (compare Figure. 3e and f of Ref. [49]) also show that the spectral weight in the sharp quasiparticle peak at the antinode vanishes between the optimally doped and underdoped UD23K compounds.

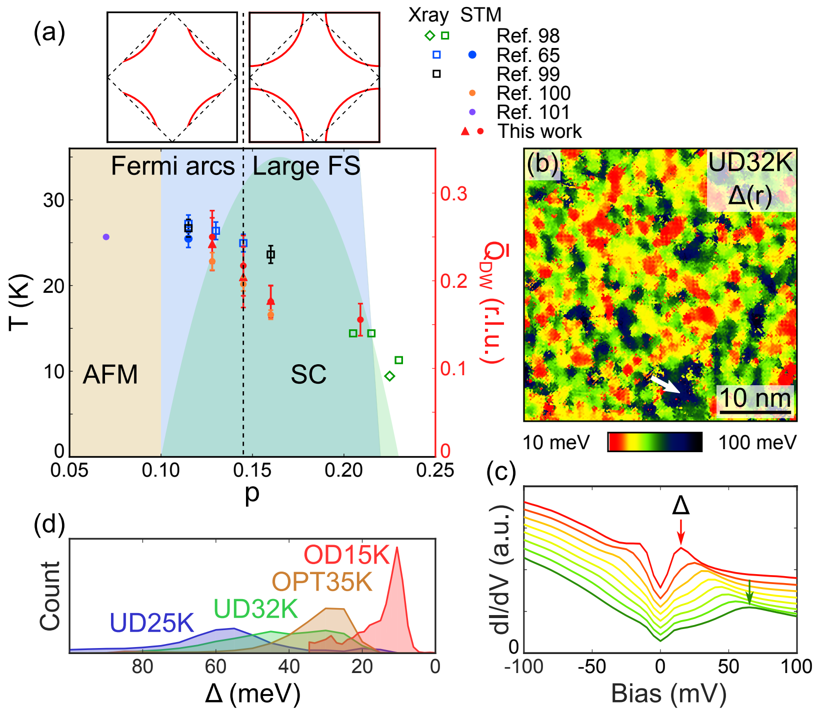

Figure 3.1: Using local electronic inhomogeneity and charge modulations to probe the Bi2201 phase diagram. **(a)** Phase diagram for Bi2201, in the region of the superconducting dome (green shading). With increasing p at low temperature, the appearance of sharp antinodal quasiparticles (insets), indicative of a large normal state Fermi surface, occurs just below optimal doping (dashed line [47]), while the spectral gap persists into overdoped compounds (blue shading, with boundary marking the closing of the gap, as measured by ARPES [50] and NMR [51]). The data points mark existing sample-average measurements of the charge modulation wavevector, $\bar{Q}_{\mathrm{DW}}$, in Bi2201 from x-ray scattering (green [98], blue [65] and black [99] open symbols) and STM (blue [65], orange [100] and purple [101] filled symbols). Red triangles and circles are the sample-average measurements of the $d-$form factor charge modulations in the x and y directions, respectively, from this work, with p determined from Ando's conversion [29], as described in the supplementary material. The green squares and diamond are the annealed and as-grown samples, respectively, from Peng et al. [98]. **(b)** The local spectral gap Δ (shown for UD32K), measured as the local minimum in the second derivative of the empty state differential conductance spectrum measured at each point. The white arrow marks the same location as in Figure 3.2h to highlight a region of large Δ contributing to the Fermi arc QPI in Figure 3.2f. **(c)** Differential conductance spectra from UD32K, averaged over spatial regions binned by Δ, and offset vertically for clarity. **(d)** Distributions of Δ within the UD25K, UD32K, OPT35K and OD15K samples. The histograms are normalized to have equal areas.

on the underdoped side. Furthermore, with the exception of Ref. [64], measurements have not taken into account the electronic inhomogeneity within these samples, even though the influence of annealing suggests that the bulk-averaged $\bar{Q}_{\mathrm{DW}}$ is sensitive to disorder [98]. Specifically, examining $\bar{Q}_{\mathrm{DW}}$ in Figure 3.1a, it is unclear (1) if a single doping-dependent incommensurate trend should be drawn through all of the measurements from $p = 0.11$ to 0.23, and (2) how the charge modulation structure in this range relates to the commensurate modulation observed in a lightly doped compound, near the insulating state [101].

Here, we use the spatial dependence of the Bi2201 density wave (DW) as a probe of the parent states in both Fermi arc and large FS regions. We find that the FS transition marks the boundary between two distinct ground states that give rise to commensurate and incommensurate charge modulations, respectively. Furthermore, concurrent mapping of the DW and FS demonstrates that conventional Fermiology is insufficient to explain the overdoped evolution of Q_{DW}.

3.1 Inhomogeneous Fermi surface transition

Within a single Bi-based cuprate crystal, an average doping of p holes per unit cell produces a highly inhomogeneous spatial distribution, resulting in large variations in the local electronic properties [70, 64, 71, 72, 73, 74, 75]. In a scanning tunneling microscope, the local electronic density of states is typically measured by the spatially-

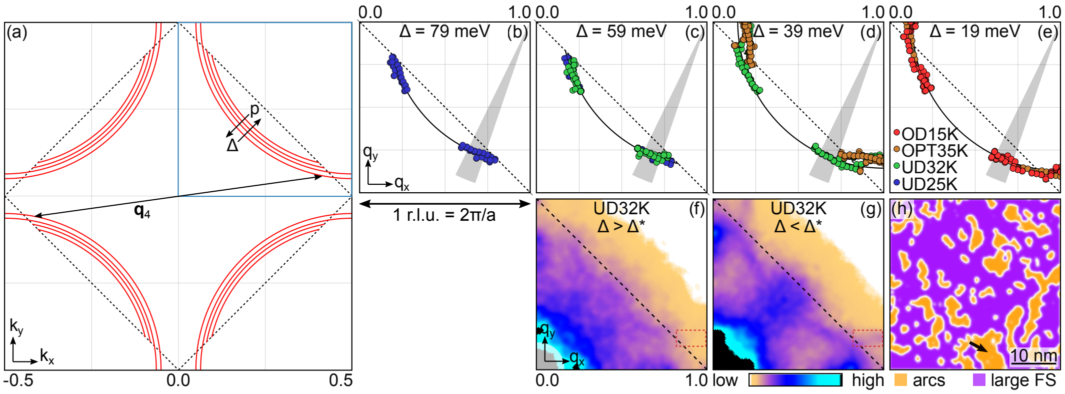

Figure 3.2: A continuous doping axis from local electronic inhomogeneity. (a) Fermi surface evolution with local doping, over the range of Δ and p studied here. The double-headed arrow indicates an example $\mathbf{q_4}$ scattering vector. The blue box marks the quadrant that is shown for the QPI panels. **(b–e)** The QPI evolution with the local gap is shown by data points marking the peak positions extracted from gap-masked $Z(\mathbf{q})$ for all four samples. Refer to Section 3.5. On each panel, data is compiled from masked regions with average Δ within a 4 meV range, centered at the indicated value. The black lines are circles determined by the average radius of the near-nodal data within a fixed angular range indicated by gray shading. For $\Delta > \Delta^*$, points are extracted only for regions inside the dashed line, as the intensity of antinodal (AN) QPI is negligible. **(f)** UD32K QPI from masking $Z(\mathbf{r})$ by the yellow regions in (h), with Δ primarily greater than Δ^*. QPI from near-nodal quasiparticles is strong but antinodal QPI (red dashed box) is not visible. **(g)** UD32K QPI from masking $Z(\mathbf{r})$ by the purple regions in (h), with Δ primarily less than Δ^*. QPI extends out to $q_{x,y} = 1$ r.l.u. (red dashed box), indicating the presence of antinodal quasiparticles arising out of the antinodal normal state Fermi surface. The gap-masked $Z(\mathbf{q})$ in (f) and (g) has been integrated from 10 meV to 25 meV, 4-fold symmetrized, and smoothed with a Gaussian filter of width 0.015 r.l.u. to reduce the appearance of noise. **(h)** Spatial division of UD32K into regions with Δ primarily greater than (yellow) or less than (purple) Δ^* (Section 3.7). The arrow in (h) marks the same location as in Figure 3.1b to highlight a region of large Δ contributing to the Fermi arc QPI. All $\mathbf{k}$ and $\mathbf{q}$ axes are in reciprocal lattice units (1 r.l.u. $= 2\pi/a$).

resolved differential conductance, $g(\mathbf{r}, E = eV) \equiv dI/dV(\mathbf{r}, V)$, where V is the sample bias, and I is the tunneling current. Binning and averaging the local spectra by gap size, $\Delta(\mathbf{r})$, as shown in Figure 3.1b and c, demonstrates the variation of the spectrum over a large field of view. Previous work has shown that smaller Δ corresponds to higher hole concentration, both locally within each sample [70, 72, 73] and globally from sample to sample [102, 103, 104, 105, 36]. Thus, the same spectrum can be found locally in samples with different global p, and the overlapping Δ distributions from the four samples studied in this work, UD25K, UD32K, OPT35K and OD15K (Figure 3.1d), allow us to move continuously from underdoped (UD) to overdoped (OD) in the phase diagram, using spatial masking [64] to hone in on a single local doping within a larger field of view (Section 3.4).

To validate the use of local doping to construct the Bi2201 phase diagram, we demonstrate that the evolution of the FS with Δ mimics that of bulk samples with p. We calculate the ratio map $Z(\mathbf{r}, E) \equiv g(\mathbf{r}, E)/g(\mathbf{r}, -E)$ [78], which enhances Bogoliubov quasiparticle interference (QPI), and eliminates artifacts associated with the tip-sample junction setup. The normal state FS can be inferred from QPI in the superconducting state as follows [106, 107, 108]. The QPI signal is dominated by wavevectors connecting regions of high density of states, *i.e.* extrema in the Bogoliubov dispersion $E(k) = \pm\sqrt{\epsilon_k^2 + |\Delta_k|^2}$, where ϵ_k is the normal state band dispersion, and Δ_k is the momentum dependent superconducting gap. For a given angle, extrema in $E(k)$ are given by $\epsilon_k = E_F$ (normal state Fermi energy), thus the Bogoliubov QPI dispersion traces out the normal state FS. In particular, it is well-established that the $\mathbf{q_4}(E)$ channel (Figure 3.2a) traces out $2k_F$ [108, 47, 48], in excellent agreement with the normal state FS measured by ARPES [108, 79]. The $\mathbf{q_4}$ wavevectors – extracted as a function of Δ by selecting a range of Δ values in $\Delta(\mathbf{r})$ to mask the $Z(\mathbf{r})$ data (Section 3.4) – together describe a single evolution of the momentum-space electronic structure extending across all samples (Figure 3.2a-e). Regions with a small gap ($\Delta = 19$ meV, Figure 3.2e) exhibit QPI tracing out a large normal state FS: the Bogoliubov quasiparticles near the antinodes at the edge of the Brillouin zone, $k_{x,y} = \pm\pi/a$ (Figure 3.2a), generate scattering with $\mathbf{q_4}$ wavevectors that extend out to $q_{x,y} = \pm 2\pi/a$. Moving to larger Δ (Figure 3.2d), the $\mathbf{q_4}$ trajectory shrinks, consistent with decreasing hole concentration, and the full evolution of the FS size, inferred from the QPI, (Figure 3.9) confirms that Δ is well-correlated to doping, in agreement with previous observations [74, 70, 64].

For larger Δ, the absence of observable QPI near $q_{x,y} = \pm 2\pi/a$ (Figure 3.3) is consistent with a normal state Fermi arc that lacks sharp antinodal (AN) quasiparticles [47, 48]. Quantitatively, the intensity of the AN QPI in the UD32K sample (Figure 3.6a) decreases with increasing Δ, before settling at a constant value, indistinguishable from the background, indicating the transition occurs at $\Delta \approx 50$ meV, which we label Δ^*. Spatially dividing the data into regions of $\Delta > \Delta^*$ and $< \Delta^*$ (Figure 3.2f, g) shows that QPI associated with a large FS and Fermi arcs both exist within the sample, such that one can move on the phase diagram from one side of the transition to the other spatially. This spatial division emphasizes that the electronic structure is determined on length scales similar to the Δ correlation length. The FS p-evolution of the cuprate phase diagram is therefore reproduced locally as a function of Δ, with the transition from Fermi arcs to large FS occurring at $\Delta^* \approx 50$ meV.

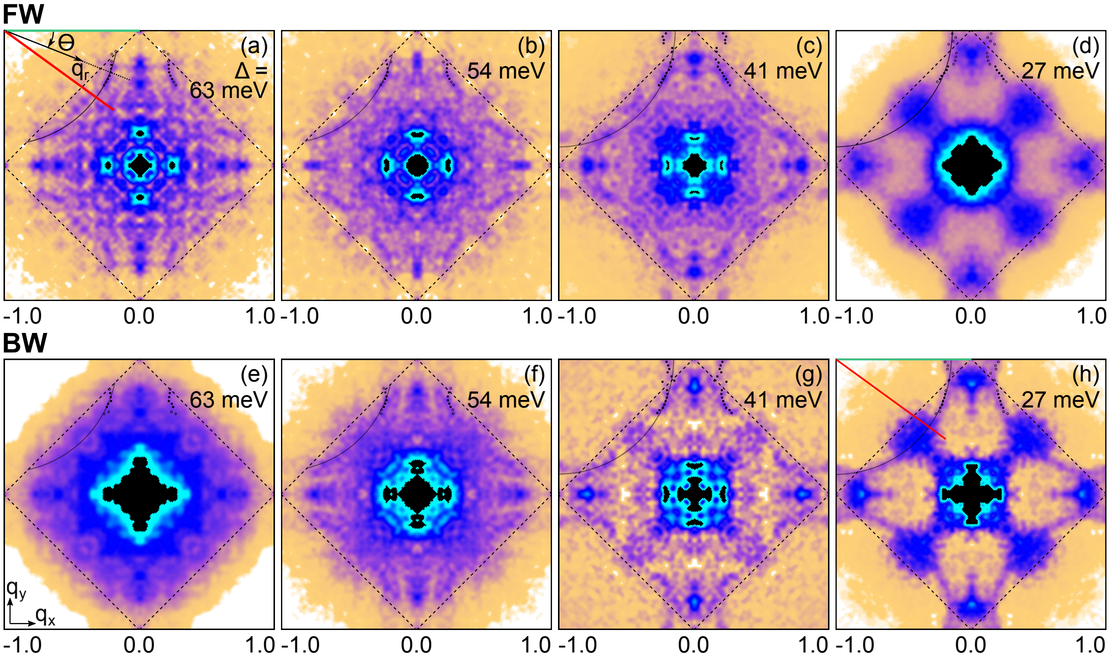

Figure 3.3: UD32K gap-masked QPI. (a–h) $Z_b^{fw}(\mathbf{q})$ (a–d) and $Z_b^{bw}(\mathbf{q})$ (e–h) are the forward and backward gap-masked $Z(\mathbf{q})$, as described in Section 3.4. The data have been integrated from 10 meV to 25 meV, four-fold symmetrized, and Gaussian smoothed with filter width 0.015 r.l.u. to reduce the appearance of noise. The evolution in intensity of AN QPI from undetectable (consistent with normal state Fermi arcs) to strong (consistent with normal state large Fermi surface, FS) is apparent in both versions of the masking, and is further analyzed in Figures 3.8 and 3.11. The points mark the locations of the $\mathbf{q_4}$ peaks, determined from cuts through the data (no smoothing), as indicated by the lines in a, where the angle of the cut, θ, and the coordinate along the cut, q_r, are defined. Axes are in reciprocal lattice units (1 r.l.u. $= 2\pi/a$). No point is shown if a peak position could not be determined reliably.

3.2 COMMENSURATE TO INCOMMENSURATE TRANSITION

We now determine the doping dependence of the charge modulation wavevector, Q_{DW}, as a function of Δ to look for signatures of the change in ground state at Δ^*. We examine $D(\mathbf{r})$, the d−form factor (dFF) component [55] of $\sum_{E>0} g(\mathbf{r}, E)/\sum_{E<0} g(\mathbf{r}, E)$, where the integration over an energy range larger than the typical Δ for each sample

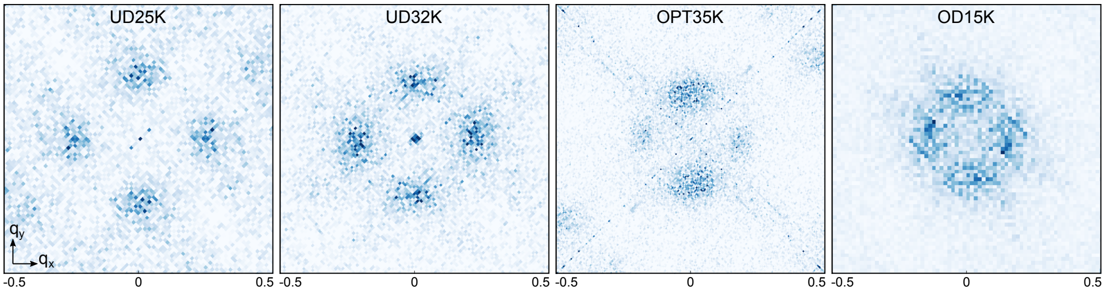

Figure 3.4: ***d*FF DW from underdoped to overdoped.** $D(\mathbf{q})$, the Fourier transform of the dFF component of $\mathcal{R}(r) = \sum_{E>0} g(r,E)/\sum_{E<0} g(r,E)$, for all samples studied in this work (no smoothing and no symmetrization). Broad but clear peaks along the q_x and q_y directions indicate the presence of dFF charge modulations in all four samples. Axes are in r.l.u.

enhances the DW signal [67]. In all four samples, the amplitude of the Fourier transform, $D(\mathbf{q})$, has broad peaks at $(\pm Q_{\mathrm{DW}}, 0)$ and $(0, \pm Q_{\mathrm{DW}})$, near the charge modulation wavevectors that have been observed by previous experiments, indicating clearly that a dFF DW exists in all (Figure 3.4). We employ the demodulation phase residue minimization technique of Mesaros *et al.* [66], to make Q_{DW} measurements robust against the strong disorder apparent from the broad shape of the DW peaks.

Comparing Figure 3.5 a–c to d–f reveals that the dFF DW exhibits distinct evolutions for small and large spectral gap regions. From 16 meV to 48 meV, Q_{DW} increases from 0.15 r.l.u. to 0.23 r.l.u., matching the change in wavevector that has been measured by resonant x-ray techniques from $p \approx 0.20$ to near-optimal doping [98, 65, 99], and consistent with an evolving incommensurate wavevector. However from $\Delta \approx 50$ meV to $\Delta \approx 85$ meV, no significant increase is observed (Figure 3.5a–c). The constant value of Q_{DW} near 0.25 r.l.u. suggests a dominant commensurate

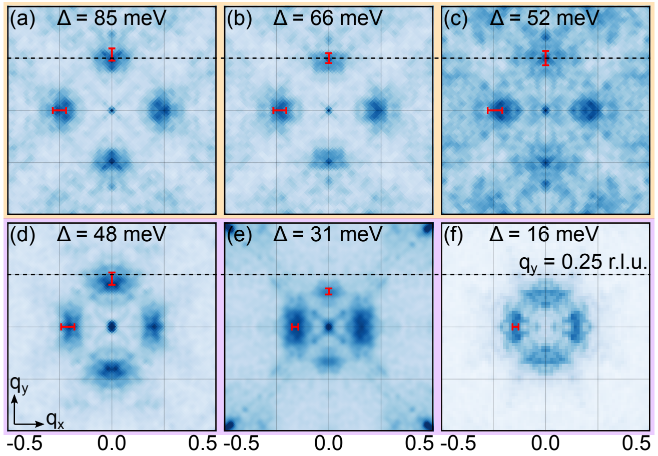

Figure 3.5: d form factor charge modulations. Examples of gap-masked $D(\mathbf{q})$ where the indicated Δ is the average value within each masked area. The data have been symmetrized along the q_x and q_y directions, and Gaussian smoothed with a 0.01 r.l.u. width. **(a–c)** Data with $\Delta > \Delta^*$ (Fermi arc regime). **(d–f)** Data with $\Delta < \Delta^*$ (large FS regime). The thin gray lines have a spacing of 0.25 r.l.u. to facilitate comparison to the Fermiology in Figure 3.2. The red symbols mark Q_{DW} determined for x and y modulation directions, where the bar length indicates ± the estimated standard deviation of spatial fluctuations within each masked area. No measurement of Q_{DW} was made for q_y modulations in OD15K (f), as the wavevector was not sufficiently well defined. The panels correspond to data from UD25K (a–c), UD32K (d), OPT35K (e) and OD15K (f).

instability. In fact, the entire $Q_{\mathrm{DW}}(\Delta)$ evolution (Figure 3.6b) is consistent with a commensurate to incommensurate transition at a location indistinguishable from Δ^*. The dashed line, showing the expected broadening of the underlying commensurate (yellow) to incommensurate (purple) trend, describes the data accurately (supplementary material Section 3.12). Furthermore, the coincident changes in Fermiology and DW commensurability strongly suggest the presence of a quantum phase transition at Δ^*.

3.2.1 COMMENSURATE DENSITY WAVE

The observed wavevector in the Fermi arc state ($\Delta > \Delta^*$) is consistent with a commensurate 4 unit cell charge modulation. The average $\bar{Q}_{\mathrm{DW}}$ from UD25K, which lies almost entirely on the underdoped side of the transition, is 0.25(3) r.l.u. and 0.24(3) r.l.u. for x and y directions, respectively, where the errors represent the estimated standard deviation of spatial variations across the entire field of view. Resonant x-ray experiments [65] have reported a doping dependent $\bar{Q}_{\mathrm{DW}}$ in this same doping range, down to $p = 0.115$. However, this apparent discrepancy, can be understood by considering: (1) the FS transition occurs near $p = 0.14$ [47] and beyond this point, we also observe an incommensurate wavevector, (2) the local inhomogeneity could plausibly induce a doping dependence of $\bar{Q}_{\mathrm{DW}}$ even in samples with average p below the transition, and (3) the strongly disordered structure of charge modulations limits the precision with which the value of Q_{DW} (and $\bar{Q}_{\mathrm{DW}}$) can be determined un-

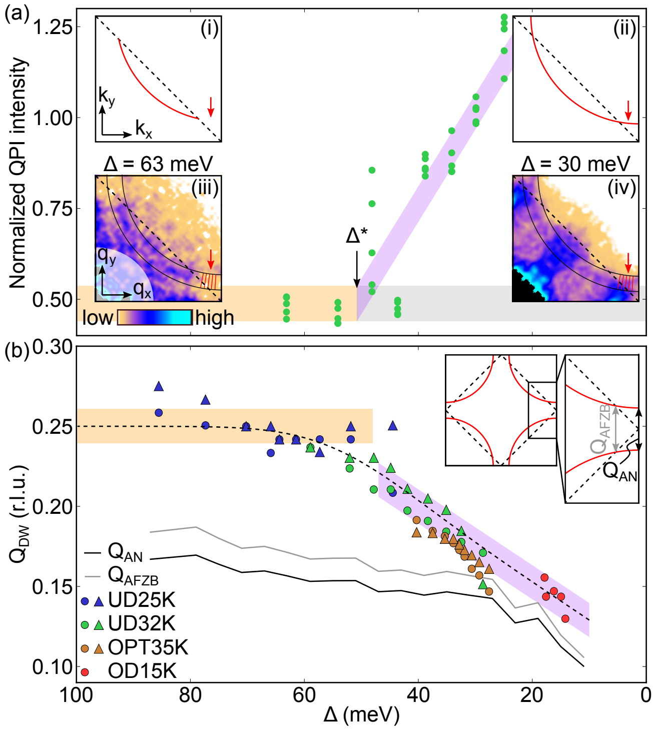

Figure 3.6: Simultaneous DW and FS transitions. (a) Intensity of AN QPI in UD32K, determined at five locations (red lines in the insets (iii, iv)), and normalized to the intensity of near-nodal QPI (Section 3.8). Δ^* indicates the approximate location where the AN QPI becomes indistinguishable from the background noise. The thick shaded lines are guides to the eye. Insets (i,ii) show the approximate Fermi surface structure inferred from two examples of the gap-masked $Z(\mathbf{q})$ (iii, iv) without and with AN QPI. **(b)** Wavevector (Q_{DW}) of the dFF charge modulations in the x (triangles) and y (circles) directions, extracted from gap-masked $D(\mathbf{r})$. Refer to Figure 3.15 for the standard deviation of Q_{DW} for each value of Δ. The thick shaded lines indicate the commensurate to incommensurate trend underlying the dashed line, which includes the expected effect of Gaussian smoothing ($\sigma = 12$ meV) in Δ, due to the resolution of the masking technique (Figure 3.7). The gray and black lines indicate the Fermiology-driven candidate wavevectors, as indicated schematically in the insets, and based on the circular Fermi surface models shown in Figure 3.2a–e.

ambiguously [66]. For the samples studied here, ignoring local doping variations, and taking the average $\bar{Q}_{\mathrm{DW}}$ value for each sample, hides the kink at the FS transition, and produces a trend of decreasing $\bar{Q}_{\mathrm{DW}}$ with doping, similar to previous reports (Figure 3.1a). Furthermore, while our data do not rule out a small doping dependence, recent experiments on Bi2212 [66, 109, 110] and very underdoped Bi2201 [101] are also consistent with a $Q = 0.25$ r.l.u. commensurate DW within the Fermi arc regime, arising from proximity to the Mott insulating state.

3.2.2 Incommensurate density wave

What is the mechanism for the incommensurate Q_{DW} observed for $\Delta < \Delta^*$? Previous work interpreted the monotonically decreasing Q_{DW} as evidence of a FS instability that follows the growing FS hole pocket [64]. In this picture, there are two natural candidates for Q_{DW}: (1) Q_{AN} that connects nested antinodal segments of the FS, and (2) Q_{AFZB} that connects the points where the FS crosses the antiferromagnetic zone boundary, the hotspots for (π, π) spin fluctuations. Both $Q_{\mathrm{AN}}(\Delta)$ and $Q_{\mathrm{AFZB}}(\Delta)$ are shown in Figure 3.6b. For the largest p (smallest Δ), Q_{DW} becomes similar to these Fermiology-derived wavevectors. However, upon decreasing p towards the transition, Q_{DW} grows more rapidly than the FS evolves. This unexpected discrepancy between Q_{DW} and Fermiology constitutes our second major finding.

3.3 Discussion

Whereas most recent theories for charge modulations in the cuprates have aimed to explain an incommensurate DW in the presence of Fermi arcs, here we are discussing (1) a *commensurate* DW in the presence of Fermi arcs and (2) an incommensurate

DW occurring in the presence of the *large FS*. This leads to two important distinctions. First, in the Fermi arc regime, strong interactions are expected, and the associated renormalization can affect a hotspot wavevector, Q_{HS}. For an instability of the large FS, however, it is not a priori clear that there should be any influence of correlations on such a Q_{HS}. This first distinction reconciles our conclusion with Ref. [65]'s interpretation that $\bar{Q}_{\mathrm{DW}}$ is Fermiology driven [‡]. We note that additional factors, such as coupling to the lattice, may affect the observed Q_{DW}. Second, theoretical studies have found that the dominant CDW instability of the large FS in the presence of exchange interactions has a wavevector along the $(\pm q, \pm q)$ direction rather than $(0, \pm q)$ or $(\pm q, 0)$, and the presence or absence of antinodal states is important in stabilizing the former or latter orientation, respectively [111]. It is therefore significant that the orientation of the charge modulations does not change at Δ^*.

To search for information about the nature of the transition at Δ^*, we consider that generically in CDW systems, a discontinuity or sharp jump in Q_{DW} occurs at a commensurate to incommensurate transition [112]. From this work, we cannot distinguish between a continuous Q_{DW} or one with a small jump, as shown by the yellow and purple trends in Figure 3.6b (see also Figure 3.7). To extract this information, Δ^* and the incommensurate $Q_{\mathrm{DW}}(\Delta)$ would need to be determined with reduced un-

---

[‡] Ref. [65] concludes that $\bar{Q}_{\mathrm{DW}}$ is explained by a FS instability, apparently in contradiction to the findings reported herein. However, Ref. [65] models an incommensurate $\bar{Q}_{\mathrm{DW}}$ in the presence of Fermi arcs, where the renormalization associated with the arc phenomenology generates a Q_{HS} significantly larger than Q_{AFZB}

certainties. However, if present, a continuous Q_{DW}, which generically requires fine tuning, would imply a mechanism-derived constraint on the position of Δ^*, not explained by existing theoretical models of the Fermi surface transition.

Ubiquity across underdoped compounds has widely been cited to motivate studying charge modulations as a route to understanding Fermi arc physics and the mechanism behind high T_c superconductivity in cuprates. However, the relationship among these three phenomena has remained an open question. The coincidence of the DW and FS transitions, observed here, establishes an intimate link between the DW and the presence of Fermi arcs, and furthermore suggests that the same interactions which generate the commensurate instability may also be responsible for the arc phenomenology.

In summary, we have reported three concrete observations: (1) there is a commensurate to incommensurate transition of Q_{DW} at a doping consistent with that of the FS transition; (2) on the underdoped side, Q_{DW} is consistent with a commensurate 4 unit cell modulation; and (3) on the overdoped side, the doping-dependence of Q_{DW} is stronger than that of the FS size.

The remainder of this chapter is dedicated to providing a more detailed description of the data processing, and demonstrates that the trends reported above can be seen directly in the data.

3.4 GAP-MASKING TECHNIQUE

The gap map $\Delta(\mathbf{r})$ is calculated by finding the position of the local minimum in the second derivative of the empty state differential conductance for the spectrum $g(\mathbf{r}, E)$ at each pixel $\mathbf{r}$, where Gaussian smoothing in energy reduces errors from noise in the data. The gap-masks are then generated by dividing the values of Δ into bins with equal counts. The mask for bin b, $M_b(\mathbf{r})$, has a value of 1 if $\Delta(\mathbf{r})$ is in b, or zero otherwise. In order to reduce periodic structure in the masks arising from the atomic corrugation or the charge modulations [64], a bilateral filter is applied to $\Delta(\mathbf{r})$ before generating the masks.

To obtain clear images in momentum transfer, $\mathbf{q}$, space, we use an additive masking technique, where for each bin b, we look for the change in Fourier transform amplitude, when adding b into the field of view. The additive masks therefore include bins summed up to a bin b: $M_{1,b}(\mathbf{r}) = \sum_{j=1}^{b} M_j(\mathbf{r})$. To avoid introducing artifacts from spatial structure of the masks, we apply Gaussian smoothing to the mask, with a spatial resolution of $1/w$. The filtered masks $\mathcal{M}_{1,b}(\mathbf{r}) = G_w(\mathbf{r}) * M_{1,b}(\mathbf{r})$, can now have any value in the range [0,1], where G_w is the Gaussian filter, and $*$ indicates convolution. This filtering imposes a spatial resolution of the masks, and implies a spread in the Δ distribution within each mask (Figure 3.7).

To explain the additive masking technique, we describe how the images in Figure 3.5 are generated. The additive masks are applied to the $D(\mathbf{r})$ map in real space:

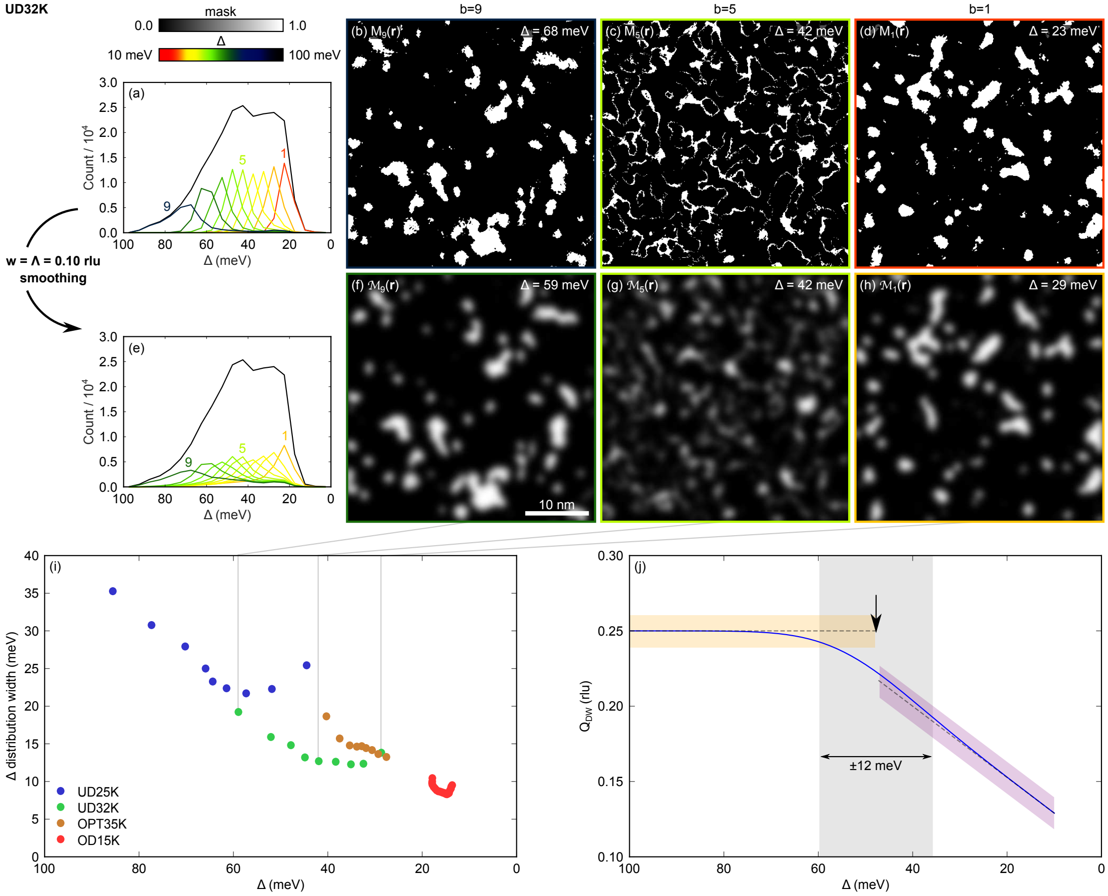

Figure 3.7: Effect of spatial resolution on $\mathbf{Q}_{\mathrm{DW}}$ measurements. The rounded shape where the overdoped and underdoped $Q_{\mathrm{DW}}(\Delta)$ trends meet is explained by the distribution of Δ within each mask. **(a)** The black curve is the histogram of Δ values in UD32K. The colored curves indicate the distributions within each mask $M_b(\mathbf{r})$. **(b–d)** Examples of masks, $M_b(\mathbf{r})$, for UD32K, generated from $\Delta(\mathbf{r})$. Even though no smoothing has been applied to the masks, a small spread in Δ results from the filtering applied to $\Delta(\mathbf{r})$. However, a spatial resolution of $1/\Lambda$ is imposed in measuring Q_{DW}. The corresponding Gaussian smoothing produces a much larger spread in Δ within each bin. **(e–h)** Same as (a–d), but with Gaussian smoothing applied to the masks, $\mathcal{M}_b(\mathbf{r})$, with smoothing parameter $w = \Lambda = 0.10$ r.l.u. **(i)** The standard deviation of Δ within the bins ($\mathcal{M}_b(\mathbf{r})$) used for the Q_{DW} analysis. **(j)** Assuming the underlying trend, $Q_{\mathrm{DW}}(\Delta)$, shown by the gray dashed line with yellow and purple shading (reproduced from Figure 4b), the blue curve shows the effect of a Gaussian distribution in Δ with standard deviation of 12 meV, similar to the distribution widths in the UD32K bins near the transition at Δ^*. Because of this smearing effect, (1) we expect to see a rounded continuous curve at a commensurate to incommensurate transition rather than a sharp kink, and (2) our Q_{DW} measurements cannot distinguish an underlying continuous meeting point of the two Q_{DW} trends from one with a small jump. The arrow indicates the location of the DW transition, and the gray shading indicates the region of ± 12 meV, where we expect the largest deviation of the data from the underlying trend.

$D_{a,b}(\mathbf{r}) = \mathcal{M}_{a,b}(\mathbf{r})D(\mathbf{r})$. To visualize the charge order peaks from bin b, we take the difference of the absolute values of the Fourier transforms, including and not including bin b. The additive masking can either add up Δ values from small to large (forward, fw) or in the reverse direction (backward, bw). For these two cases, the gap-masked $D(\mathbf{q})$ is

$$\begin{aligned} D_b^{fw}(\mathbf{q}) &= D_{1,b}(\mathbf{q}) - D_{1,b-1}(\mathbf{q}) && (3.1) \\ D_b^{bw}(\mathbf{q}) &= D_{b,N}(\mathbf{q}) - D_{b+1,N}(\mathbf{q}), && (3.2) \end{aligned}$$

where N is the total number of bins and $D_{a,b}(\mathbf{q})$ is a real number, the amplitude of the Fourier transform. Unless otherwise specified, functions of $\mathbf{q}$ refer to amplitudes of Fourier transforms. The panels of Figure 3.5 have Gaussian smoothing applied with width 0.01 r.l.u., to reduce the appearance of noise.

The QPI shown in the insets of Figure 3.6, are $Z_b^{fw}(\mathbf{q})$, generated from $Z(\mathbf{r})$ using this same technique. Gap-masked QPI from UD32K in Figure 3.3 shows the evolution of the Fermi surface structure across the full range of Δ within the sample. These QPI images have Gaussian smoothing applied with width 0.015 r.l.u., and have been four-fold symmetrized.

For analysis of the Fermi surface structure, we used overlapping bins, *i.e.*

$$Z_b^{fw}(\mathbf{q}) = Z_{1,b}(\mathbf{q}) - Z_{1,b-n}(\mathbf{q}) \tag{3.3}$$

$$Z_b^{bw}(\mathbf{q}) = Z_{b,N}(\mathbf{q}) - Z_{b+n,N}(\mathbf{q}), \tag{3.4}$$

in order to measure the QPI wavevectors at finer spaced intervals of Δ, with more variations in the masks. The integer n determines the overlap between subsequent bins.

The average value of the gap for bin b is then:

$$\Delta_b^{fw} = \frac{\sum_{\mathbf{r}} \Delta(\mathbf{r}) \left[\mathcal{M}_{1,b}(\mathbf{r}) - \mathcal{M}_{1,b-n}(\mathbf{r})\right]}{\sum_{\mathbf{r}} \left[\mathcal{M}_{1,b}(\mathbf{r}) - \mathcal{M}_{1,b-n}(\mathbf{r})\right]}, \tag{3.5}$$

$$\Delta_b^{bw} = \frac{\sum_{\mathbf{r}} \Delta(\mathbf{r}) \left[\mathcal{M}_{b,N}(\mathbf{r}) - \mathcal{M}_{b+n,N}(\mathbf{r})\right]}{\sum_{\mathbf{r}} \left[\mathcal{M}_{b,N}(\mathbf{r}) - \mathcal{M}_{b+n,N}(\mathbf{r})\right]}, \tag{3.6}$$

where $n = 1$ for distinct, as opposed to overlapping, bins.

Here, we used $N = 19$ for UD25K, 31 for UD32K, 31 for OPT35K and 21 for OD15K, with $n = 4$, and $w = 0.20$ r.l.u.

For $D(\mathbf{q})$ in Figure 3.5, we used $N = 9$ with $n = 1$ for UD25K, UD32K and OPT35K, and $N = 21$ with $n = 4$ for OD15K. For OD15K, only the data for bins $b = 5$, 9, 13, 17, and 21 are plotted in Figure 3.6b, to present data from distinct bins, as were used on the other 3 samples. The mask smoothing parameter, w, was chosen to include the resolution, Λ, of determining Q_{DW} (Section 3.9): $w = \Lambda = 0.10$ r.l.u.

for UD25K, $w = \Lambda = 0.10$ r.l.u. for UD32K, $w = \Lambda = 0.04$ r.l.u. for OPT35K and $w = \left(\Lambda^{-2} + 0.20\,\text{r.l.u.}^{-2}\right)^{-1/2}$ with $\Lambda = 0.04$ r.l.u. for OD15K.

3.5 Fermi surface structure from gap-masked QPI

$Z(\mathbf{q})$ was integrated over the low energy layers to capture the full dispersion of the Bogoliubov quasiparticles, from 1.5 meV to 9 meV for OD15K, 5 meV to 5 meV for OPT35K, 10 meV to 25 meV for UD32K, and 5 meV to 15 meV for UD25K.

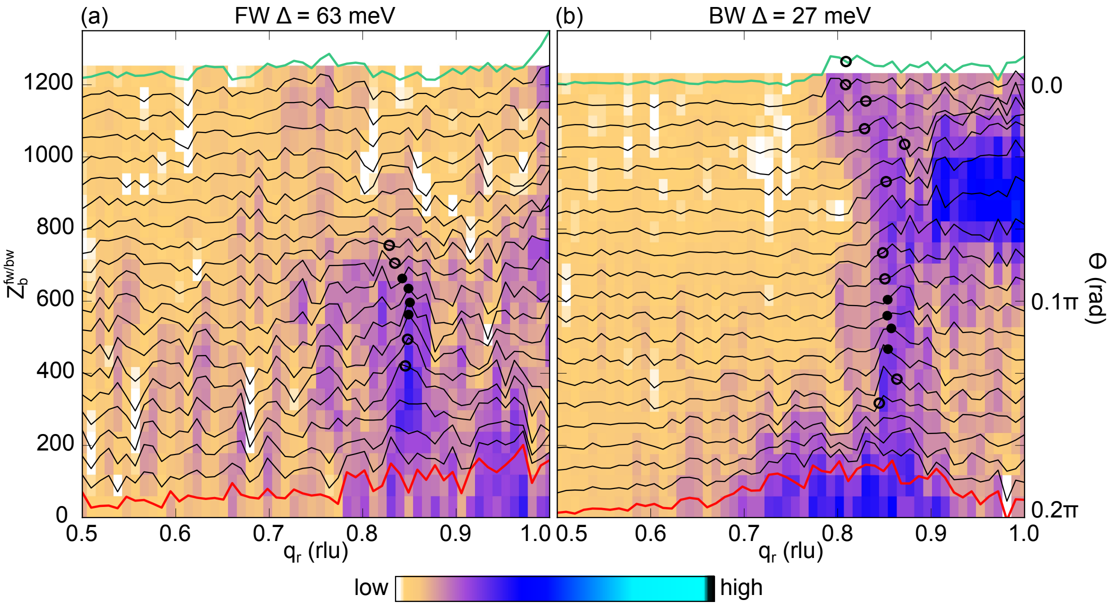

Figure 3.8: UD32K gap-masked QPI. The linecuts for (a) and (h) of Figure 3.3 are shown in (a) and (b), respectively, by the waterfall plots and the colorscale. The cuts have been averaged over a transverse width of 8 pixels (0.074 r.l.u.). Cuts are offset vertically (by increments of 60) for clarity. The points mark the positions of the $\mathbf{q_4}$ QPI peaks. The $\mathbf{q_4}$ data points, thus determined, for all samples, are compiled as a function of Δ to describe the Fermi surface evolution in Figure 3.2 (b–e). Data points used to determine the Fermi surface radius in Figure 3.9 are indicated by filled (as opposed to open) symbols. No point is shown if a peak position could not be determined reliably.

For each Δ bin, the QPI wavevectors, as shown in Figure 3.2, were extracted from the positions of peaks in one-dimensional cuts through $Z_b^{fw}(\mathbf{q})$ and $Z_b^{bw}(\mathbf{q})$, and are shown in Figures 3.3 and 3.8 for UD32K. To quantitatively determine the size of the Fermi surface, the wavevectors from all samples were binned together, as shown in Figure 3.2. A circular hole pocket was determined from the average radius of the data in the range $\theta = 0.105\pi$ to 0.145π, where θ is defined in Figure 3.3, and the range was selected because this near-nodal QPI is consistently measured across the Fermi surface transition (unlike the antinodal QPI), and is least influenced by nearby scattering channels or the DW signal. The evolution of the Fermi surface radius is shown in Figure 3.9, and explicitly demonstrates that Δ tracks the local doping.

3.6 CONSISTENCY ACROSS SAMPLES

Using Δ as a local doping axis for the cuprate phase diagram relies on two requirements:

1. Δ is monotonically related to doping.
2. On the lengthscale of the analysis (*i.e.* the size of patches within the masks), the same electronic structure is found in regions with the same average Δ, independent of the sample.

The first requirement is supported sample to sample by measurements across multiple cuprate families [36, 113]. STM work has also demonstrated a correlation of Δ to the local concentration of dopants [70, 72], and the relationship of FS size to Δ in

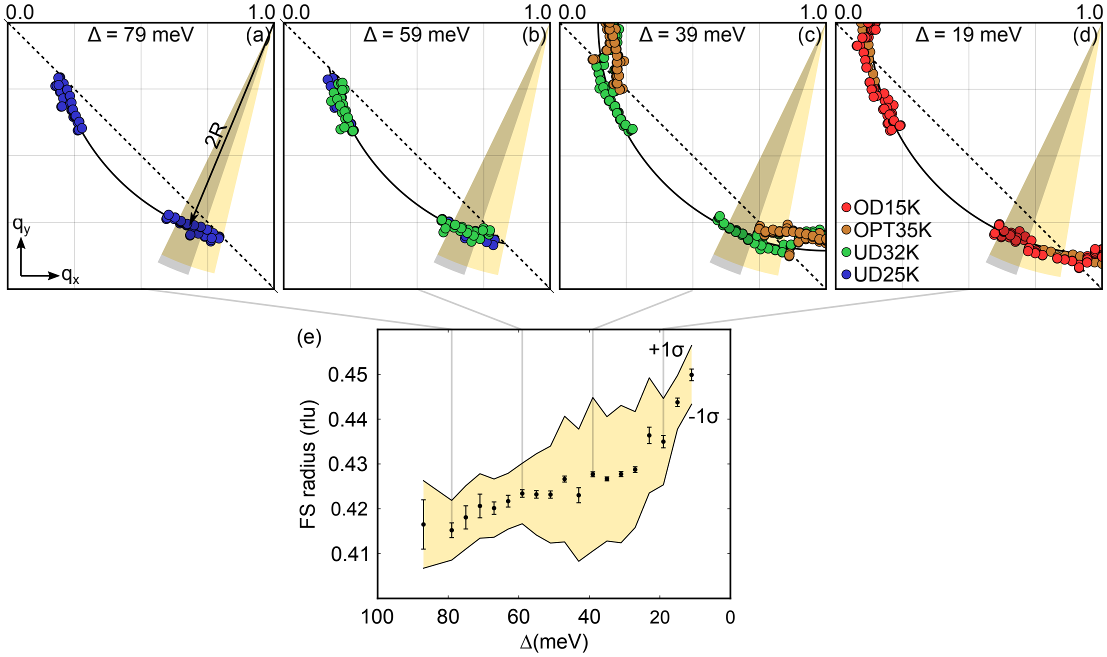

Figure 3.9: Quantifying the Fermi surface size. (a–d) $\mathbf{q_4}$ QPI data, as shown in Figure 3.2. The data within the fixed angular range indicated by the gray shading is used to estimate the radius of the normal state Fermi surface (R). **(e)** Normal state Fermi surface radius inferred as half the measured $\mathbf{q_4}$ QPI dispersion radius. The error bars indicate standard error, as a lower bound error estimate. As an upper bound error estimate, we also show the standard deviation (with respect to R) for a larger range of data, corresponding to the Fermi surface structure up to the AFZB ($\pm 1\sigma$ lines and yellow shading). The increase in Fermi surface size with decreasing Δ validates that Δ is a proxy for local doping. This circular model of the Fermi surface is used to determine Q_{AN} and Q_{AFZB} shown in Figure 3.6.

Figure 3.9 concretely confirms the local doping effect on the band structure. The consistency across samples is demonstrated by the extracted $\mathbf{q_4}$ wavevectors in Figure 3.2 and by the Q_{DW} data points in Figure 3.6.

Figure 3.10 demonstrates a remarkable consistency between samples directly in the gap-masked Fourier transformed data. The QPI from the $\Delta = 27.3$ meV bin of UD32K matches the QPI from the OPT35K bin with the same size Δ. Both are described well by the dashed red circle with radius twice the FS radius from Figure 3.9. Moving to the $\Delta = 62.1$ meV bin of UD32K, the QPI is dramatically changed, but agrees well with the equivalent bin from UD25K.

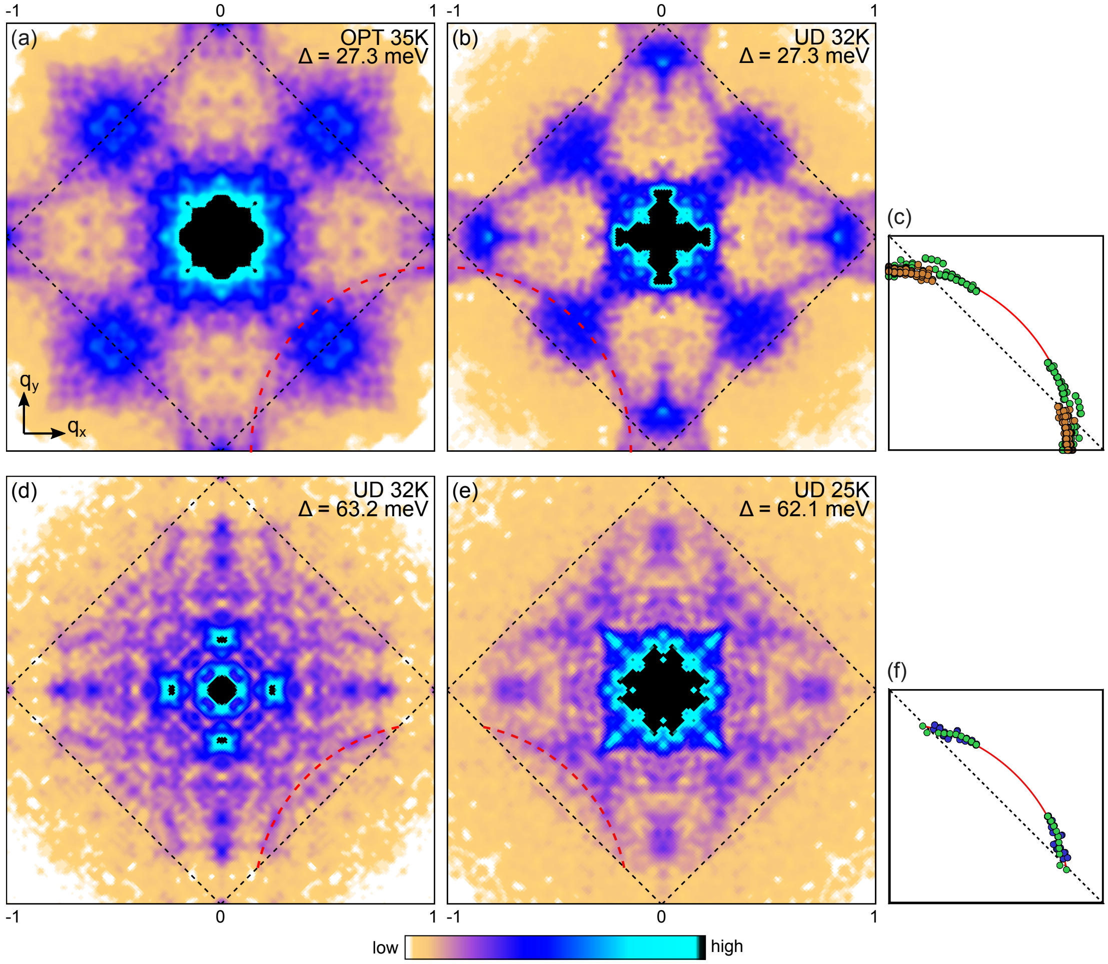

Figure 3.10: Gap-masked QPI, comparison across samples: Gap-masked QPI, $Z_b^{fw/bw}(\mathbf{q})$ for $\Delta \approx 27$ meV **(a,b)** and 63 meV **(d,e)**. The red curves are circular arcs with a radius twice the inferred FS radius, as determined in Figure 3.9. The same arc is shown in all panels for each of the two Δ values. **(c,f)** $\mathbf{q_4}$ QPI data points compiled from all four samples from all masked regions with average Δ in the ranges 25 meV< $\Delta_b^{fw/bw}$ <29 meV (c) or 61 meV< $\Delta_b^{fw/bw}$ <65 meV (f). Consistent QPI is observed across samples for regions with the same size Δ. Axes are in r.l.u.

3.7 FERMI ARC AND LARGE FERMI SURFACE QPI

To generate Figure 3.2f and g, we locate the bin b^*, which has the largest average gap value below 49 meV, the average of Δ^* estimates based on forward and backward masked data shown in Figure 3.6a and Figure 3.9b. Bins $b \leq b^*$ primarily have $\Delta < \Delta^*$, and $b > b^*$ primarily $\Delta > \Delta^*$. This division is only approximate due to the Gaussian smoothing of the masks (Figure 3.7). To probe QPI from the two regimes divided by Δ^*, panels f and g show $Z_{b^*+1,N}(\mathbf{q})$, and $Z_{1,b^*}(\mathbf{q})$, respectively. Note that in this case, there is no subtraction after applying the Fourier transform. Figure 3.2h shows $\mathcal{M}_{1,b^*}(\mathbf{r})$, where the colorscale interpolates between 0.0 (yellow) and 1.0 (purple). Because the masks used to obtain f and g are related by $\mathcal{M}_{b^*+1,N}(\mathbf{q}) = 1 - \mathcal{M}_{1,b^*}(\mathbf{q})$, the yellow and purple indicate the regions primarily contributing to the Fermi arc and large Fermi surface QPI, respectively.

3.8 LOCATING THE FERMI SURFACE TRANSITION IN UD32K

In order to quantify the disappearance of the antinodal (AN) QPI, which is apparent directly in the data (Figure 3.3), we measure the intensity at cuts spaced at regular angular intervals, as shown in Figure 3.11. The cuts are averaged over a transverse width of 0.07 r.l.u., and a length determined by the QPI radius (twice the FS radius) ± 0.06 r.l.u. To compare QPI intensities from different bins, the intensities of the cuts are normalized by the average intensity of five cuts closer to the nodal QPI (dashed red lines in Figure 3.11a). Figure 3.6a tracks this normalized intensity for five cuts near the antinode (red lines in Figure 3.6a-iii and -iv).

To confirm the trend of vanishing AN QPI from small to large Δ, Figure 3.11b tracks the intensity of AN QPI for the backward-masked data (equivalent to Figure 3.6a). The data display the same trend, independent of the masking scheme: from small Δ to large Δ, the intensity of the AN QPI decreases, and then flattens out where it is indistinguishable from the background noise.

To further confirm that the observed behavior is associated with the FS transition, where AN QPI should disappear but nodal QPI should remain intense, we track the angular dependence of the QPI intensity in (c) and (d). There are two different behaviors: (1) for large Δ, the QPI intensity drops off significantly from $\pi/8$ to $\pi/16$, meaning that the AN QPI is weak, as expected to arise from Fermi arcs in the normal state, and (2) for small Δ, the QPI intensity decreases only slightly from $\pi/8$ to 0, indicating strong AN QPI, as expected to arise from a large Fermi surface in the normal state. The Fermi surface transition is therefore quantified by the intensity of the AN QPI.

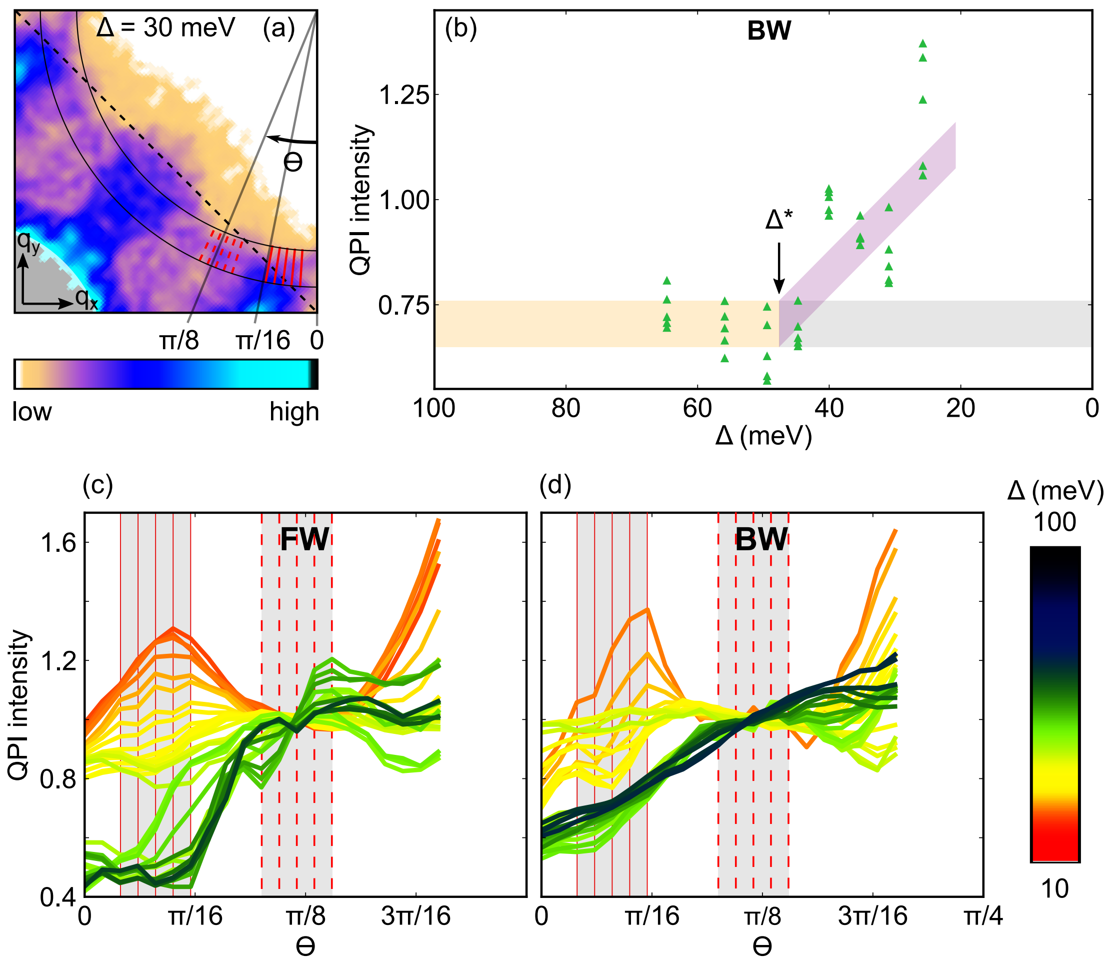

Figure 3.11: Locating the Fermi surface transition in UD32K. **(a)** $Z_b^{fw}(\mathbf{q})$, which has been integrated from 10 meV to 25 meV, four-fold symmetrized, and Gaussian smoothed with filter width 0.015 r.l.u. The black lines indicate the length over which the cuts are integrated, *i.e.* the QPI radius (Figure 3.9) ± 0.06 r.l.u. Figure 3.6a and panel b show the intensity from the five cuts indicated as solid red lines, for the forward-masked data and for the backward masked data, respectively. θ is the angle of the cut. **(b)** Intensity of antinodal QPI in the backward-masked UD32K data. The intensity has been normalized by the average of 5 cuts near $\theta = \pi/8$, indicated by the red dashed lines in (a,c,d). **(c,d)** The angular dependence of the normalized QPI intensity, as a function of θ for $Z_b^{fw}(\mathbf{q})$ and $Z_b^{bw}(\mathbf{q})$, respectively. Each trace represents a single bin, b, with the colorscale indicating the bin's average Δ.

3.9 EXTRACTING $\mathbf{Q}_{\mathrm{DW}}$ FROM dFF CHARGE MODULATION

We follow the procedure described by Mesaros [66], where Q_{DW} is determined as the wavevector which minimizes the demodulation residue, $R_{\mathbf{Q}} = \sqrt{|R^x_{\mathbf{Q}}|^2 + |R^y_{\mathbf{Q}}|^2}$, over the field of view. In the case of a strongly disordered density wave, this measurement has a more clearly defined interpretation than fitting peaks of the Fourier transformed data. Both techniques are compared in Section 3.11. A detailed explanation of demodulation residue can be found in Reference [66]. In this section, we extend the technique for application to masked regions of the data's field of view.

$D(\mathbf{r})$ is demodulated by the reference wavevector $\mathbf{Q}$ in

$$\tilde{\Psi}_{\mathbf{Q}}(\mathbf{q}) = \exp\left(\frac{-q^2}{2\Lambda^2}\right)\tilde{\psi}(\mathbf{q}+\mathbf{Q}), \tag{3.7}$$

where $\tilde{\psi}(\mathbf{q}) = \tilde{D}(\mathbf{q})$ over a domain that isolates the charge order peak, $\tilde{D}$ is the complex-valued Fourier transform of $D(\mathbf{r})$, and $\tilde{\Psi}_{\mathbf{Q}}$, $\tilde{\psi}$ are complex-valued functions. The Gaussian cutoff imposes a spatial resolution of $1/\Lambda$.

The Q_{DW} measurement proceeds in each gap bin, b, by integrating the residue only over the masked region:

$$\begin{aligned} R^{\alpha}_{b,\mathbf{Q}}[\Psi] &= \int d^2\mathbf{r}\,\mathcal{M}_b(\mathbf{r})\ \mathrm{Re}\left[\Psi^*_{\mathbf{Q}}(-i\partial_\alpha)\Psi_{\mathbf{Q}}\right], \\ \sigma^2 &= \frac{\int d^2\mathbf{r}\,\mathcal{M}_b(\mathbf{r})\ \sum_\alpha \left[\Psi^*_{\mathbf{Q}}(-i\partial_\alpha)\Psi_{\mathbf{Q}}\right]^2}{\int d^2\mathbf{r}\,\mathcal{M}_b(\mathbf{r})\ |\Psi_{\mathbf{Q}}|^2}, \end{aligned} \tag{3.8}$$

where α is either x or y, $\Psi_{\mathbf{Q}}(\mathbf{r})$ is the inverse Fourier transform of $\tilde{\Psi}_{\mathbf{Q}}(\mathbf{q})$, and with $\mathbf{Q} = \mathbf{Q}_{\mathrm{DW}}$, σ estimates the standard deviation of spatial fluctuations in the modulation wavevector.

For this analysis, we used $w = 0.0$ as the mask smoothing parameter for UD25K, UD32K and OPT35K, with $\Lambda = 0.10$ r.l.u., 0.10 r.l.u. and 0.04 r.l.u., respectively. For OD15K, we used $w = 0.20$ r.l.u. and $\Lambda = 0.04$ r.l.u. The samples with smaller wavevector require a smaller Λ for Q_{DW} to be robust against the choice of Λ.

3.10 VISUALIZING THE RELATIONSHIP BETWEEN Q_{DW} AND THE FERMI SURFACE

In this section, we demonstrate that the discrepancy between Q_{DW} and Fermi surface derived wavevectors (Figure 3.6) is apparent directly from the gap-masked Fourier transformed data.

Figure 3.12 shows the QPI and the *d*FF DW peaks with the data on the same scale such that the wavevectors can be compared directly. Because $\mathbf{q}_4 = 2\mathbf{k}_F$, where $\mathbf{k}_F$ are the normal state Fermi wavevectors, the vector from the Bragg peak to the AN QPI (blue arrow) is Q_{AN}. However, The wavevectors of the AN QPI cannot be inferred directly from the signal intensity (blue circle) at $q_{x,y} = 1.0$ r.l.u. The *d*FF DW, although low in intensity at the energies of the QPI dispersion, appears as satellite peaks around the Bragg peaks. Thus the peak in intensity near $q_{x,y} = 1.0$ r.l.u. is shifted away from Q_{AN}, causing the data points in this region to be shifted to larger

q from the location of the QPI. This interpretation is supported by the signal's bias voltage dependence: this signal moves outward from the energy of the Bogoliubov QPI to the energy of maximum dFF DW intensity. Thus Q_{AN} is estimated by extrapolating the QPI from the region farther from the antinode which is not affected by the DW signal. Q_{AFZB} can be seen directly from the location where the QPI crosses the dashed line. The DW peaks are located at larger $\mathbf{Q}$ than the AN QPI, demonstrating the discrepancy with a simple FS-derived instability.

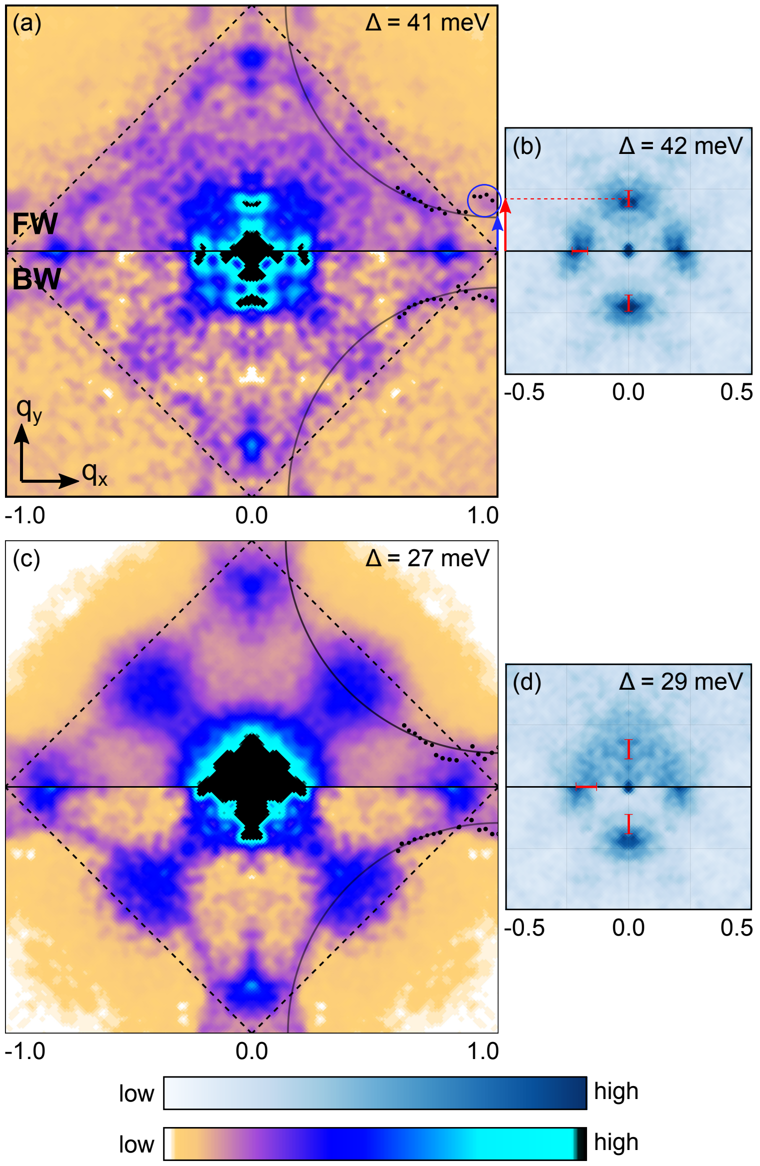

Figure 3.12: Comparison of QPI to *d*FF DW peaks in UD32K. (a,c) Gap-masked QPI ($Z_b^{fw}(\mathbf{q})$ and $Z_b^{bw}(\mathbf{q})$). The points mark specific QPI wavevectors, as described in Section 3.5. Q_{AN} is indicated by the blue arrow. The top (bottom) half of each panel is the forward (backward) version of the gap-masking. **(b,d)** Gap-masked $D(\mathbf{q})$, the Fourier transform of the *d*FF component of $\mathcal{R}(r) = \sum_{E>0} g(r,E)/\sum_{E<0} g(r,E)$, with peaks located at $(\pm Q_{\mathrm{DW}}, 0)$ and $(0, \pm Q_{\mathrm{DW}})$. The red arrow indicates Q_{DW}. While forward and backward gap-masked $D(\mathbf{q})$ are shown, Q_{DW}, determined by phase residue minimization, (red markers) is independent of masking direction, and identical markers are therefore duplicated on both top and bottom halves of each panel.

3.11 Comparison of Q_{DW} measurement techniques

The demodulation residue minimization technique of determining the density wavevector is simpler to understand in the absence of amplitude modulations and in one dimension (Figure 3.13). In this case, the wavevector of a modulation $\psi(r) = \cos[\phi(r)]$ minimizes the accumulated phase offset between ψ and the reference modulation $\cos(Qr)$ [66], as shown schematically by the spatial dependence of the modulation phase ϕ (black curve in a) and the demodulation residue R_Q for $Q = Q_1$, Q_2, and Q_3.

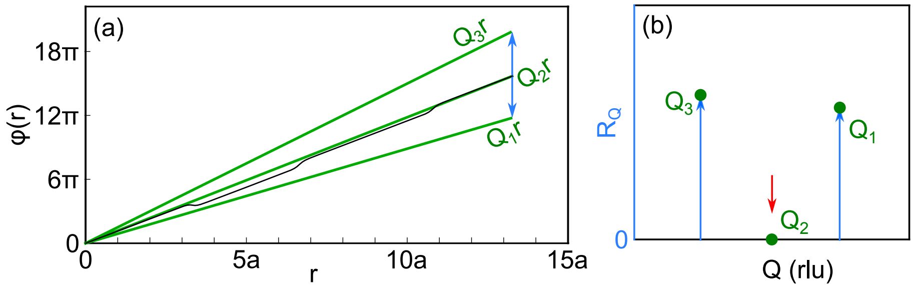

Figure 3.13: Demodulation residue minimization in 1D (schematic). (a) Spatial (r) dependence of the phase ϕ for a phase-disordered modulation $\psi(r) = \cos[\phi(r)]$ (black), and for reference modulations $\cos(Qr)$ with $Q = Q_1$, Q_2, and Q_3. **(b)** Demodulation residue R_Q for $Q = Q_1$, Q_2, and Q_3. In this example, we have assumed a uniform amplitude for the modulation in ψ.

The modulation wavevector can also be determined by fitting the d-form factor (dFF) density wave (DW) peaks in the Fourier transformed data. Isolating the dFF components in the gap-masked D (Figure 3.14a), the dFF DW appears as four peaks around (0,0). A Gaussian peak is fitted to a linecut through one of the peaks to determine the peak location. Figure 3.15 compares Q_{DW} from both techniques. The two measurements are in reasonable agreement.

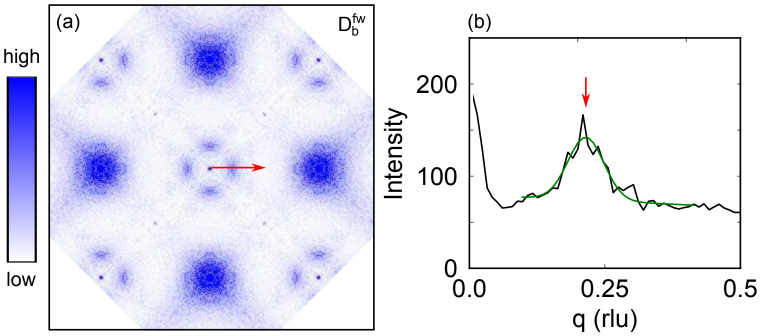

Figure 3.14: Peak fitting (UD32K). (a) Example *d*FF component of the gap-masked D. The *d*FF DW appears as four peaks around (0,0). The data has been four-fold symmetrized. No smoothing has been applied. **(b)** Example of a linecut through D_b^{fw} for UD32K (black), with the fitted peak shape (green).

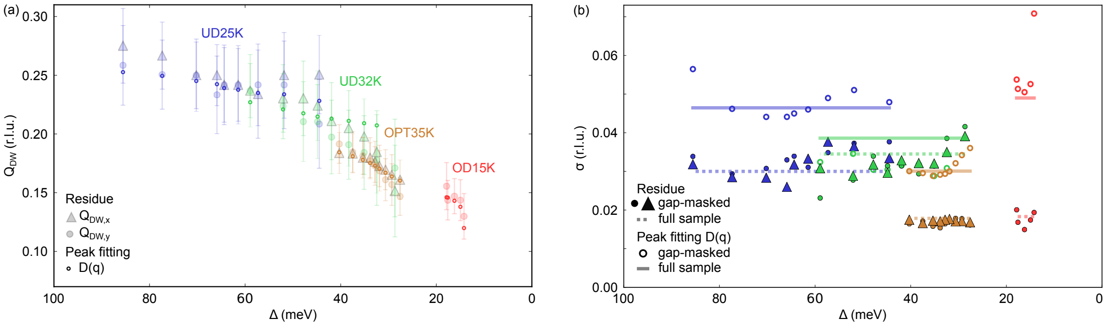

Figure 3.15: Comparison of $\mathbf{Q}_{\mathrm{DW}}$ measurement techniques. (a) Q_{DW} determined from Gaussian fitting to cuts through $D_b(\mathbf{q})$. Each point shown is the average of the values measured from forward and backward masked data. The peak fitting analysis was done using the exact same Δ binning as for the demodulation phase residue minimization, shown in Figure 3.6, and reproduced here as the faded symbols, with errorbars corresponding to the estimated standard deviation of spatial fluctuations within each masked region. The peak-fitting measurements are in agreement with the phase residue minimization Q values, but in the case of strong disorder in the modulation structure, such as is observed in Bi-2201, Q_{DW} from demodulation phase residue minimization has a clearer interpretation [66]. **(b)** Standard deviation of the Gaussian fits to the DW peaks from each gap-masked region (open circles), and from the un-masked data (solid lines). The peak width for the masked data is influenced by disorder in the DW structure as well as the masking procedure. For comparison, we also show the estimated standard deviation of spatial fluctuations in Q_{DW} estimated in the phase residue technique (same as the error bars in a), both for the gap-masked regions and for each entire sample. We expect that systematic errors in these σ values, associated with the data processing and fitting could be as high as ~0.02 r.l.u. The sensitivity of the electronic inhomogeneity to the details of the sample annealing [98, 74] may contribute to the variation in peak width across samples.

3.12 Commensurate-incommensurate transition model

In this section, we address in detail the consistency of our data with a commensurate to incommensurate density wave (DW) transition at Δ^*, defined as the location of the FS transition.

First, we characterize the change in slope of $Q_{\mathrm{DW}}(\Delta)$ from the overdoped to the underdoped regime. The data shows a bend which occurs near Δ^*, and as described in Figure 3.7, we expect within $\sim$ 12 meV of this point, the data will be an average of information from the states on both sides of the transition. To quantify the change in slope observed, we perform linear regressions on either side, excluding a 12 meV region centered at Δ^* (49 meV is the average of the Δ^* estimates derived from the decreasing antinodal QPI intensities in forward, Figure 3.6a, and backward, Figure 3.11b, masked data):

$$Q_{\mathrm{DW}} = m_{\mathrm{arc}}\Delta + b_{\mathrm{arc}}, \qquad \Delta > \Delta^* + 6 \text{ meV} \tag{3.9}$$

$$Q_{\mathrm{DW}} = m_{\mathrm{lrg}}\Delta + b_{\mathrm{lrg}}, \qquad \Delta < \Delta^* - 6 \text{ meV} \tag{3.10}$$

The resulting estimates are $m_{\mathrm{lrg}} = 0.0022(2)$ r.l.u./meV and $m_{\mathrm{arc}} = 0.0011(2)$ r.l.u./meV, shown in Figure 3.16(a). The errors reported are the statistical errors associated with the regression. Naively, m_{arc} would seem to imply inconsistency with a commensurate wavevector in the Fermi arc state; however, we expect systematic errors larger than the statistical errors may result from proximity to the transition, as our data does not extend to large enough Δ to produce a reliable fit if we exclude a large region of data around Δ^*, *i.e.* the fit includes a range of data that we expect may be affected by our finite Δ resolution. It is also possible that the nature of defects in the DW structure may be related to the small observed slope for $\Delta > \Delta^*$, but investigation of this structure is beyond the scope of this work.

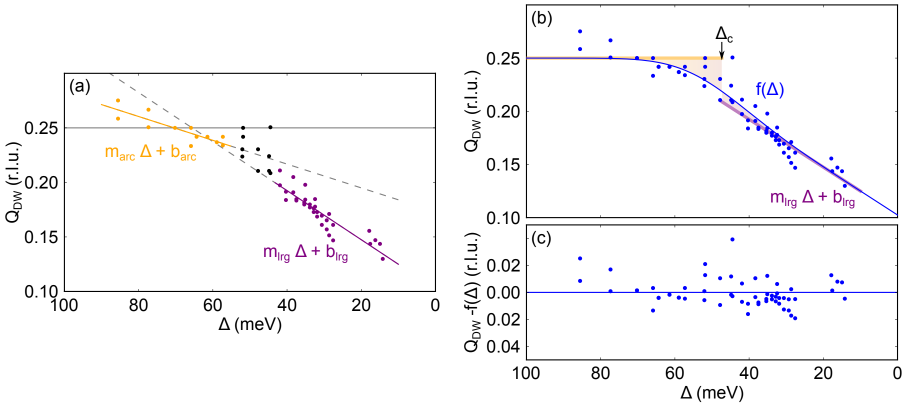

Figure 3.16: Density wave transition (a) *d*-form factor DW wavevector measurements (Q_{DW} reproduced from Figure 3.6b) with linear regressions in the Fermi arc (yellow) and large FS (purple) regimes, excluding the data (black points) in a 12 meV range centered at the FS transition, $\Delta^* \approx 49$ meV. **(b)** Q_{DW} measurements [same as in (a)], fitted to $f(\Delta)$, described in the accompanying text. The shaded yellow and purple lines show the commensurate to incommensurate Q_{DW} trend, $f_0(\Delta)$. **(c)** Residuals from fitting Q_{DW} to $f(\Delta)$.

In Figure 3.16b, c, we demonstrate that despite the nonzero slope of $Q_{\mathrm{DW}}(\Delta)$ for $\Delta > \Delta^*$, the data are consistent with a commensurate Q_{DW}. We build a model based on a commensurate $Q_{\mathrm{DW}} = 0.25$ r.l.u. for $\Delta > \Delta_c$, and an incommensurate Q_{DW}, for $\Delta < \Delta_c$, from the linear regression in Equation 3.10. We then add a Gaussian convolution in Δ to account for the finite Δ-resolution of our measurements:

$$\begin{aligned} f_0(\Delta) &= \begin{cases} 0.25 \text{ r.l.u.} & \Delta > \Delta_c \\ m_{\mathrm{lrg}}\,\Delta + b_{\mathrm{lrg}} & \Delta < \Delta_c \end{cases}, && (3.11) \\ G(\Delta) &= \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(\frac{-\Delta^2}{2\sigma^2}\right), && (3.12) \\ f(\Delta) &= G * f_0. && (3.13) \end{aligned}$$

The model is then fitted to the data by optimizing Δ_c, the location of the transition, only, as m_{lrg} and b_{lrg} are fixed from the initial linear regression, and σ is fixed at 12 meV, based on the minimum Δ distribution width in UD32K, which we expect to be relevant near the transition. The resulting model $f(\Delta)$ describes the data well, and is shown in Figure 3.16b, with residuals in c. The fitted value of Δ_c, 47 meV, is in good agreement with Δ^*.

From this analysis, we therefore can conclude that: (1) within the experiment's resolution, the change in slope is located at Δ^*; (2) the finite Δ resolution of our measurements explains both the broadened, rather than abrupt, change and the finite slope observed for $\Delta > \Delta^*$, although further study would be required to definitively rule out a small slope in this regime; (3) the Q_{DW} measurements are consistent with a commensurate Q_{DW} for $\Delta > \Delta^*$.

4

Local correlations in disordered materials with neural networks

Inhomogeneity is widespread in materials of interest. Strongly correlated materials have a tendency toward phase segregation [7]. Symmetry broken phases can form complex domain structures. Chemical substitution and doping (leading to local variations in chemical environments) are common strategies for designing new materials and tuning their properties.[*] In disordered materials, the wealth of information in spatially resolved measurements encodes physical relationships underlying the material properties. The spatial location is a tuning parameter, and the correlations between observables can be accessed by disentangling variations. As has been pointed out by others, the technology for measuring structural and electronic properties has advanced to a point where large multidimensional data sets can be acquired in short periods of time, making large spatially resolved data sets accessible but also making the process of extracting information from the data a crucial challenge motivating the application of machine learning (ML) [109, 116, 117].

ML, now widespread in technology and rapidly permeating into scientific research, excels at analyzing large data sets, modeling multidimensional data in large multidimensional parameter spaces, and functioning robustly in the presence of variations such as noise. These strengths are points of failure for standard techniques of analyzing the electronic structure of inhomogeneous materials, where we typically introduce multiple stages of analysis to reduce data down to a manageable set from which we can fit a low dimensional model[†]. This process does not always take full advantage of

---

[*] To name a few, Fe- and Cu-based high temperature superconductors are obtained through doping, the quantum anomalous hall effect was first first realized by doping magnetic atoms into $(Bi,Sb)_2Te_3$ to generate the required ferromagnetism [114], and substituting Te for Se in FeSe induces a band crossing reported to bring on a topological superconducting state [115].

[†] In Chapter 3, we reduced large spatial maps containing many thousands of data points acquired via STM down to 18 points per sample.

the information in the data set and can impose assumptions at each stage of data reduction. Furthermore, a common strategy to understand spatial variations is to consider relatively homogeneous patches individually, e.g. by masking, but analyzing small regions can reduce signal to noise. The ability of ML to empirically model complex structures and thus test hypotheses on the full data set may relieve many limitations of standard data analysis techniques. While many types of ML may be useful in solving these problems, here, we consider the application of supervised learning with convolutional neural networks (CNN) to uncover the physical model underlying spatial inhomogeneity in a cuprate superconductor.

Obtaining good training data is crucial to successful NN models. A neural network (NN) is a versatile function (as will be described in Section 4.1), where information is "learned" empirically from a large data set during the training process.[‡] The desired NN outputs have to be known for the training data, and typically large data sets are required, as diversity within the set prevents the NN from overfitting and enables the NN to ignore noise and disorder when making predictions on new data. For these reasons, it is often not feasible in condensed matter physics to use real experimental data for training. Several recent works using neural networks to analyze experimental data have trained on synthetic data [109, 117, 118]. Generating synthetic data sets

---

[‡] This can be thought of in contrast to choosing a simple functional model (based on physical knowledge) to represent data, and then performing regression, where information is encoded both in the choice of the model function and during regression. While physical knowledge can be integrated into NN architectures (and the choice of training data) [109], the restrictions are typically much less severe than in typical regression analyses.

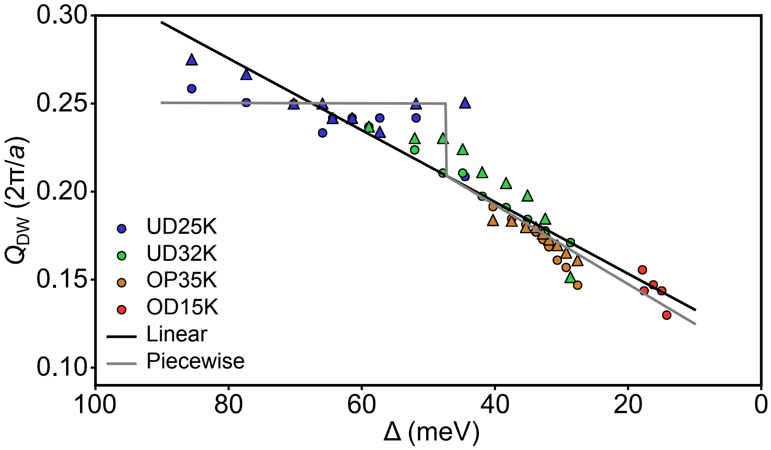

Figure 4.1: $Q_{\mathrm{DW}}(\Delta)$ **hypotheses.** The linear and piecewise hypotheses (lines) representing a single incommensurate instability and a commensurate to incommensurate transition, respectively, were determined by fitting the Q_{DW} data measured in Chapter 3 (circles and triangles) to $Q_{\mathrm{DW}} = m_l\Delta + b_l$ and to a piecewise function with a transition from $Q_{\mathrm{DW}} = 0.25$ r.l.u. to $Q_{\mathrm{DW}} = m_p\Delta + b_p$ at Δ_c. The parameterizations are $m_l = 2.0$ r.l.u./eV, $b_l = 0.113$ r.l.u., $m_p = 2.2$ r.l.u./eV, $b_p = 0.102$ r.l.u., and $\Delta_c = 47$ meV. Fitting for the piecewise hypothesis is described in Section 3.12.

that mimic the real data is one of the central challenges in applying NNs to analyze experimental data.

In this chapter, we describe the application of CNNs to address the question: Are the disordered *d*FF charge modulations observed in Bi2201 governed by a single incommensurate doping-dependent wavevector $Q_{\mathrm{DW}}(\Delta)$, or is there a commensurate to incommensurate transition? In either case, the $Q_{\mathrm{DW}}(\Delta)$ relationship should produce the Q_{DW} measurements from Chapter 3, so we make our question precise by fitting two models, linear (incommensurate) and piecewise (commensurate to incommensurate), to the data (Figure 4.1). The goal is to build a CNN that classifies a pair of spatially resolved data—$\mathcal{R}(\mathbf{r})$ and $\Delta(\mathbf{r})$—according to the underlying $Q_{\mathrm{DW}}(\Delta)$

relationship (Figure 4.2). While NNs have been used to analyze material structure from atomic resolution TEM data [117], to classify quantum gas microscope data according to the theoretical model underlying disordered antiferromagnetic correlations [118], and to measure the sample average dFF DW wavevector in Bi2212 [109], to our knowledge, NNs have yet to be applied to analyze the correlation between multiple experimentally measured observables in the electronic structure of quantum materials. In this work, we demonstrate that CNNs can be trained to classify pairs of Δ and synthetic dFF DW images according to the underlying $Q_{\mathrm{DW}}(\Delta)$ relationship, laying the groundwork for applying these types of measurements to real experimental data.

This chapter is organized as follows. We begin with a brief introduction to NNs in Section 4.1 and then motivate the crucial role of the training data in shaping the NN model in Section 4.2. Section 4.3 analyzes the structure of the dFF DW by enumerating and quantifying the forms of disorder in $\mathcal{R}(\mathbf{r})$ such that the data can be mimicked with stochastic variations in synthetic training data. In Section 4.4, we describe data simulation, evaluate the synthetic data in comparison to the real data, and demonstrate that the synthetic data encodes the $Q_{\mathrm{DW}}(\Delta)$ hypothesis. In Section 4.5, we demonstrate training CNNs to classify synthetic data. In Section 4.6, we discuss implications for extending measurements to experimental data and for using ML to analyze local correlations in general. Finally, in Section 4.7, we propose a strategy for building statistical variations into the hypotheses, in order to emphasize qualitative

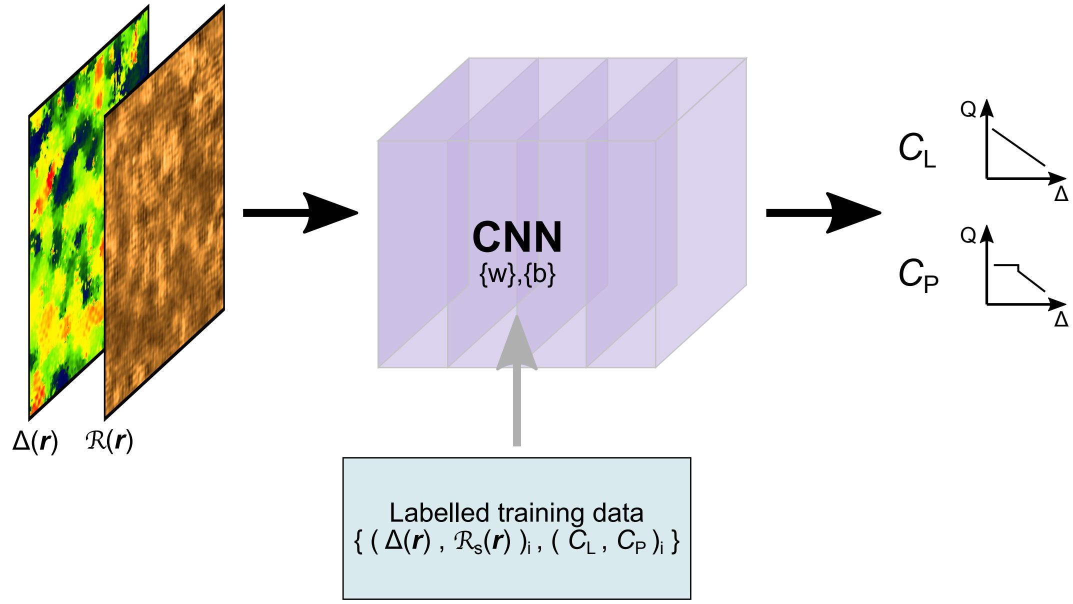

Figure 4.2: Schematic of $Q_{\mathrm{DW}}(\Delta)$ classification using a CNN. A pair of $\Delta(\mathbf{r})$ and $\mathcal{R}(\mathbf{r})$ images are the inputs to a convolutional neural network, whose output (C_L, C_P) classifies the input data as corresponding to either the linear $(C_L > 0.5, C_P < 0.5)$ or piecewise $(C_L < 0.5, C_P > 0.5)$ $Q_{\mathrm{DW}}(\Delta)$ hypothesis. The CNN model parametrized by its neurons' weights $\{w\}$ and biases $\{b\}$, is determined by the set of training data [labelled by the desired CNN outputs, where the training labels $(C_L, C_P)_i$ are binary, $(0, 1)$ or $(1, 0)$] during the training process. The weights and biases are also influenced by other choices for the training process.

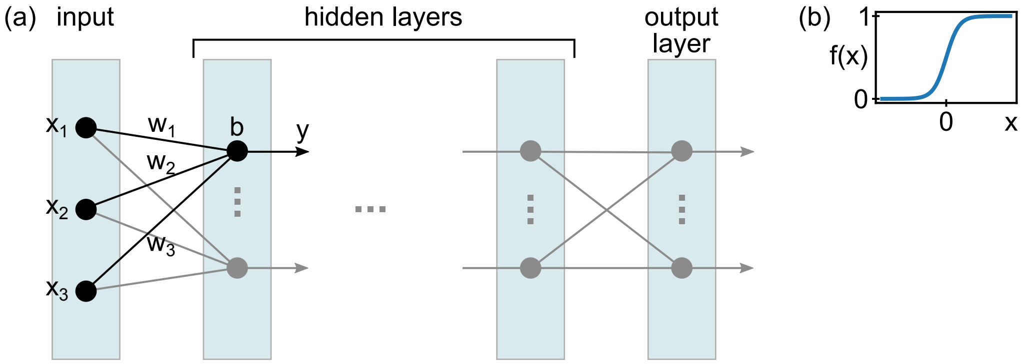

Figure 4.3: **(a)** Schematic of a fully connected neural network. In black, the action of one neuron is highlighted. **(b)** Sigmoid activation function.

rather than quantitative distinctions.

4.1 BRIEF INTRODUCTION TO NEURAL NETWORKS

Here, we give a brief overview of neural networks (NNs). For more details, there are excellent textbooks[§].

The neuron is the functional unit of a NN (Figure 4.3). It calculates a linear combination of inputs $x_1, \ldots, x_N$ and then, crucially, applies a non-linear activation function σ to produce the output y:

$$y = \sigma\left(\sum_{i=1}^{N} w_i x_i + b\right) \tag{4.1}$$

The neuron's parameters – the weights w_i and the bias b – are optimized during train-

---

[§] For an introduction, refer to Nielsen [119], and for a more in depth text, refer to Goodfellow *et al.* [120]

ing. The activation function (sigmoid, softmax, hyperbolic tangent and rectified linear unit are some of the most common) allows the neuron's output to effectively turn on or off depending on whether or not the weighted sum of inputs crosses the threshold set by the bias, and the weights determine which inputs are most important in this "decision."

Assembling neurons into a series of layers generates a highly versatile function[¶], again parametrized by the set of all the constituent weights and biases. In a fully connected layer (also called "dense"), each neuron takes input from all neurons in the previous layer. Convolutional layers, standard in NNs used for image recognition, convolve (formally the operation is a cross-correlation rather than a convolution [120]) the previous layer with a filter and then optionally apply an activation function. Typically, a layer applies multiple filters (not sequentially, but independently) in order to produce a series of feature maps. Each filter detects a pattern (or feature) such that the resulting convolution indicates where the feature is located in the input data. Applying multiple convolutional layers builds the complexity of the features that can be detected and increases nonlinearity (if activation is applied).

[¶] Fully connected NNs, even with just a single layer, are able to approximate any continuous function [121, 122]. This property is referred to as universality.

4.2 ROLE OF TRAINING DATA

Application of neural networks (NNs) can be understood as regression, interpolation and extrapolation. The training process is the regression step where we optimize the NN parameters to describe the mapping from the training data instances to their labels. Predicting on unlabeled data (in our case, experimentally measured STM data) is using the optimized model to interpolate or extrapolate. The training data – both the information contained within each instance and how these instances are distributed in the input space – shape the resulting model.

In this work, we primarily look for two qualities in our training data:

1. Training data should be similar to the real data
2. Training data should contain representative variation

In this section, we motivate these goals for our synthetic training data.

To understand the role of labeled training data in determining the model "learned" by the NN, recall the process of linear regression. The goal is to find the function $f(x) = mx + b$, which is most likely to describe data set (x_i, y_i), with i ranging from 1 to N. Given a choice of parameters m and b, one can define the probability of measuring a data point (x_i, y_i), and conversely the set of data points defines a likelihood associate with a given choice of m and b:

$$L = \prod_i \frac{1}{\sqrt{2\pi}\sigma_i} \exp\left[\frac{-(y_i - f(x_i))^2}{2\sigma_i^2}\right], \tag{4.2}$$

where we assume that each y_i is sampled from a normal distribution with standard deviation σ_i. We typically minimize $\sum_i (y_i - f(x_i))^2/2\sigma_i^2$ over m and b in order to maximize the likelihood. In this process, our initial measured data (x_i, y_i) has defined a function that we use to make predictions (interpolation and extrapolation).

The neural network functions similarly. Our training data are analogous to x_i data, but are multidimensional. Each x_i corresponds to a pair of 256 by 256 pixel maps of $\Delta(\mathbf{r})$ and synthetic $\mathcal{R}(\mathbf{r})$. The labels are analogous to y_i and are also multidimensional, consisting of a confidence score for each of our $Q_{\mathrm{DW}}(\Delta)$ hypotheses. Finally, our function f_{NN} is determined by the architecture and the weights and biases of our NN. (While f in linear regression is parametrized by m and b, NNs typically contain a far greater number of parameters.) A number of different "loss functions" can be used for optimization in place of the sum over squared deviations, and the typical metrics used to gage success of the training process are classification accuracy on the training data set and on an independent labeled "validation" data set.

We now return to goal 1 for our training data. In the linear regression example, we expect smaller errors in interpolated and extrapolated y values in the domain near our measured data. We expect larger errors in extrapolating to values of x far from the data range. The same applies to training NNs. We want our training data to span a similar range of input space to our real data, *i.e.* in qualitative terms, our synthetic training data should look like our STM data.

Finally, our data contains many forms of disorder and noise. By including representative disorder in the synthetic data, we train the NN to ignore these unimportant variations in the input space, hence goal 2.

4.3 DW STRUCTURE IN REAL SPACE

To meet the defined goals for our synthetic data, in this section, we examine the structure of the *d*FF charge modulations in $\mathcal{R}(\mathbf{r})$ with the goal of uncovering features that we aim to reproduce in the synthetic data.

For use with neural networks, we standardize the pixel spacing of our four data sets to 5 px/a (pixels per unit cell), and we crop the data [$\Delta(\mathbf{r})$ and $\mathcal{R}(\mathbf{r})$] into uniformly sized instances of data each 256 by 256 px ($51.2a$). Figure 4.4 shows an instance of data for each of the four samples, as well as the Fourier transform $\mathcal{R}(\mathbf{q})$.

For this chapter, we expand on the notation introduced in Chapter 2 to describe the DW structure. A summary of notation is provided in the glossary at the beginning of this dissertation. To simplify notation and to make directionality explicit, we define $\mathbf{Q}_x$ and $\mathbf{Q}_y$ to be the average *d*FF DW wavevectors ($\bar{Q}_{\mathrm{DW}}$) for x- and y-direction modulations, respectively. Unless otherwise specified, $\Psi_{\mathbf{Q}}$ is the demodulated *d*FF component of $\mathcal{R}$, *i.e.* we use $\tilde{\psi}(\mathbf{q}) = \tilde{D}(\mathbf{q})$ in Equation 2.10.

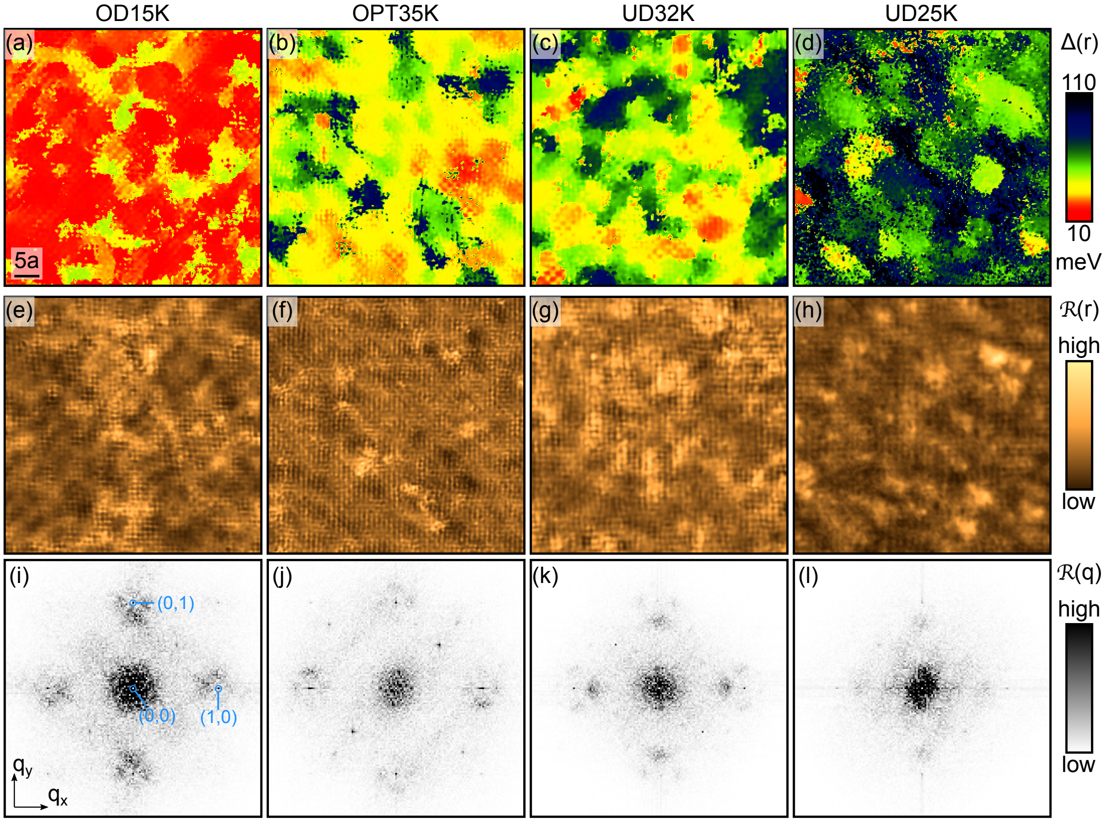

Figure 4.4: Real space DW structure in four samples. (a–d) Spatially varying spectral gap $\Delta(\mathbf{r})$ for one instance of data for each sample. **(e–h)** $\mathcal{R}(\mathbf{r})$ over the same fields of view as in a–d. Disordered ladder structures associated with *d* form factor modulations are visible, along with many forms of disorder. **(i–l)** Fourier transforms of e–h. The *d*FF modulations appear as four broad peaks around the Bragg wavevectors, although these peaks may also contain weight from an *s* component. Coordinates are labeled in r.l.u.

4.3.1 $Q_{\mathrm{DW}}(\Delta)$ INHOMOGENEITY

As demonstrated in Chapter 3, the local doping inhomogeneity [imaged as $\Delta(\mathbf{r})$] creates a spatially varying wavevector $Q_{\mathrm{DW}}(\mathbf{r})$ (Figure 3.6b). For this work, we therefore assume that the locally favored wavevector, *i.e.* the ideal (i) inhomogeneous wavevector, can be expressed as a function of the locally measured spectral gap:

$$Q_{\mathrm{DW}}^{(i)}(\mathbf{r}) = h[\Delta(\mathbf{r})]. \tag{4.3}$$

We expect data to include deviations from $Q_{\mathrm{DW}}^{(i)}(\mathbf{r})$ as a result of strain and other types of disorder, as will be discussed in the following sections. Classifying the function h according to the linear and piecewise hypotheses of Figure 4.1 is the ultimate goal of this work.

4.3.2 DIRECTIONALITY

From the four peaks in $D(\mathbf{q})$ shown in Figure 3.4, it is clear that all samples display both x- and y-direction modulations. Due to the presence of multiple signals beyond the dFF charge order and the short correlation length, how the two directions of modulation are related locally is not clear from visual inspection of $\mathcal{R}(\mathbf{r})$. The dFF modulations in Bi2212 form nanoscale unidirectional domains [67, 66, 123], and resonant x-ray scattering experiments suggest that charge order in YBCO is similarly locally

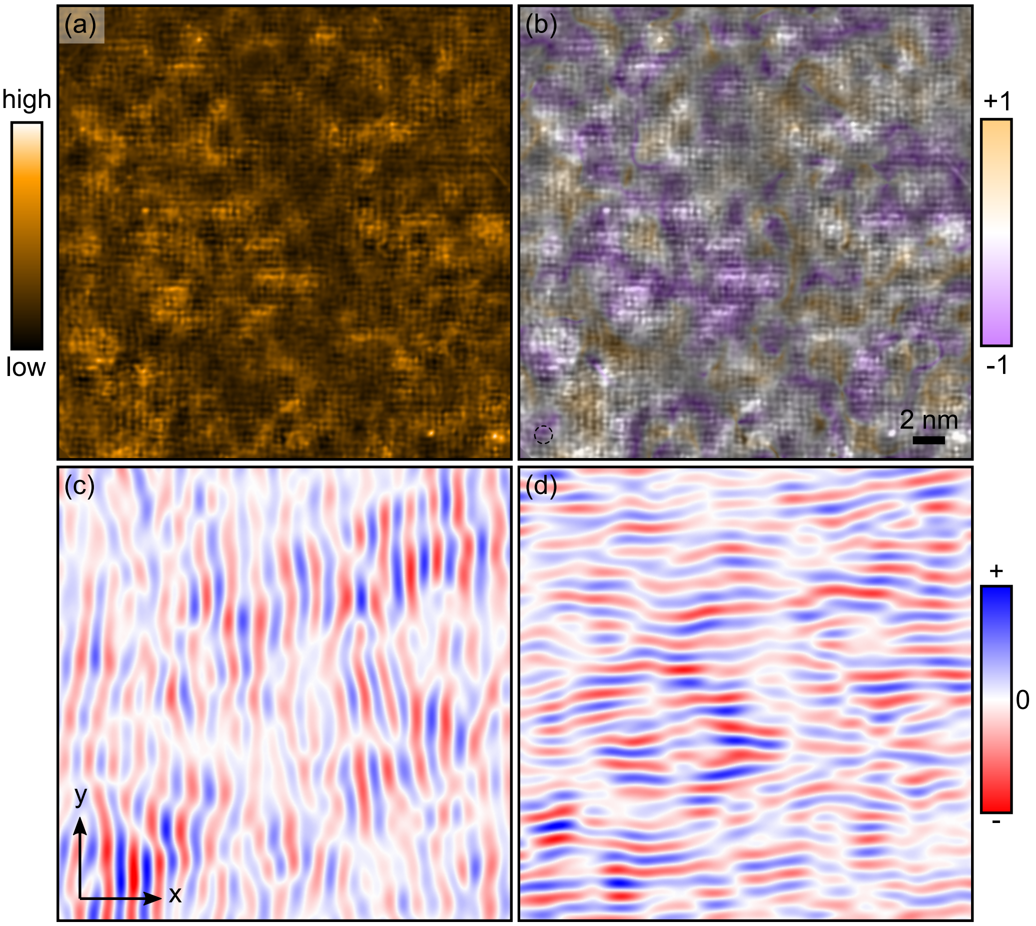

Figure 4.5: DW directionality in UD32K. (a) $\mathcal{R}(\mathbf{r})$. **(b)** dFF DW domain structure in $F(\mathbf{r})$ overlaid on $\mathcal{R}(\mathbf{r})$. The orange (purple) regions have stronger $x(y)$-direction modulations. Regions dominated by x- and y-direction modulations are intertwined on short lengthscales. **(c)** $\mathrm{Re}[D_x(\mathbf{r})]$ showing the x-direction dFF modulations. **(d)** $\mathrm{Re}[D_y(\mathbf{r})]$ showing the y-direction dFF modulations. The dashed circle in (b) has radius Λ^{-1} to indicate the spatial resolution used to extract the x- and y-direction modulations.

unidirectional [124]. In Figure 4.5, we extract the x- and y-direction modulations in $D_x(\mathbf{r})$ and $D_y(\mathbf{r})$ defined in Fourier space:

$$\tilde{D}_\alpha(\mathbf{q}) = \exp\left[-\frac{|\mathbf{q}-\mathbf{Q}_\alpha|^2}{2\Lambda^2}\right]\tilde{D}(\mathbf{q}), \tag{4.4}$$

where $\alpha = x, y$, and $\mathbf{Q}_\alpha$ is the average dFF DW wavevector ($\bar{Q}_{\mathrm{DW}}$) for α-direction modulations. Note that like $\tilde{\Psi}_{\mathbf{Q}_\alpha}(\mathbf{q})$ from Equation 2.10, $\tilde{D}_\alpha(\mathbf{q})$ isolates a single charge order peak, but because the peak remains at $\mathbf{Q}_\alpha$, $D_\alpha(\mathbf{r})$ displays the modulation instead of being demodulated. In addition, because we keep only the peak at $\mathbf{Q}_\alpha$ and not the peak at $-\mathbf{Q}_\alpha$, $D_\alpha(\mathbf{r})$ is complex and the modulation appears in $\mathrm{Re}[D_\alpha(\mathbf{r})]$. In both $D_x(\mathbf{r})$ and $D_y(\mathbf{r})$, nanoscale patches of higher and lower amplitude are visible.

To look for domain structure, we calculate the asymmetry [67]:

$$F(\mathbf{r}) = \frac{|\Psi_{\mathbf{Q}_x}(\mathbf{r})| - |\Psi_{\mathbf{Q}_y}(\mathbf{r})|}{|\Psi_{\mathbf{Q}_x}(\mathbf{r})| + |\Psi_{\mathbf{Q}_y}(\mathbf{r})|}. \tag{4.5}$$

While $F(\mathbf{r})$ displays patches with a dominant directionality, there are also regions where the amplitudes are comparable and $F(\mathbf{r})$ varies rapidly. The intertwining of x- and y-direction modulations on smaller length scales than in Bi2212 agrees with the difference in correlation lengths between the two materials.

4.3.3 Correlated amplitude and phase variations

As described in Section 2.4, the amplitude and phase of $\Psi_{\mathbf{Q}_\alpha}(\mathbf{r}) = A_{\mathbf{Q}_\alpha}(\mathbf{r})\exp[i\Phi_{\mathbf{Q}_\alpha}(\mathbf{r})]$ approximate the DW amplitude and phase, where $\Phi_{\mathbf{Q}_\alpha}(\mathbf{r})$ indicates a phase offset from reference modulation $\cos(\mathbf{Q}_\alpha \cdot \mathbf{r})$.

The dFF DW (Figure 4.6 examines the y-direction modulations from one region of UD32K) exhibits highly correlated amplitude and phase variations. The phase has one dimensional regions of high gradient $|\nabla\Phi_{\mathbf{Q}_\alpha}|$ corresponding to regions of strongly suppressed amplitude (*e.g.* the red dashed curve). Previous work has shown that in Bi2212, the local phase of s/s'FF charge modulations is correlated to defect position [70]. The structure we observe here might be expected to arise if material defects locally pin the DW phase, leading to high phase gradients at the boundaries between pinned regions. Discommensurations arising from a competition between lattice commensurate and incommensurate instabilities [84] could also produce boundaries of rapidly changing phase. In the presence of this type of domain-like disorder, it is logical that the amplitude should be suppressed in regions where the modulation phase gradient differs strongly from the favored $Q_{\mathrm{DW}}^{(i)}$. It is also possible that the local directionality could be related to the observed structure.[‖] Empirically, tracking the DW amplitude as a function of the locally varying phase gradient (Figure 4.7) explicitly

---

[‖] Mesaros *et al.* have correlated the positions of topological defects (which appear preferentially in low amplitude regions) to boundaries between regions of dominant x- and y-directionality [123].

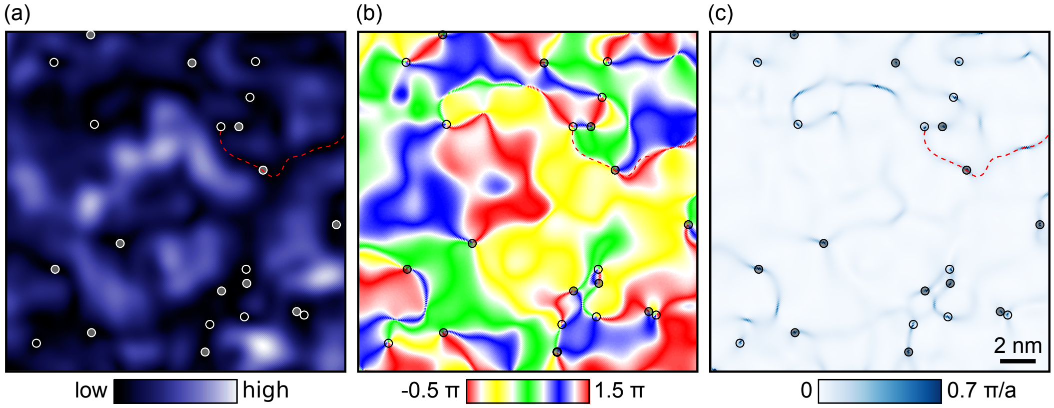

Figure 4.6: dFF DW defects in UD32K. (a) $A_{\mathbf{Q}_y}$, the amplitude of the demodulated y-direction dFF DW $\Psi_{\mathbf{Q}_y}(\mathbf{r})$, for the same instance of UD32K data shown in Figure 4.4. **(b)** $\Phi_{\mathbf{Q}_y}(\mathbf{r})$, the complex phase of the demodulated y-direction dFF DW $\Psi_{\mathbf{Q}_y}(\mathbf{r})$. Topological point defects around which the phase winds by $\pm 2\pi$ are seen as points where all four colors meet, and are marked in all panels by circles. The ordering of the colors on a path enclosing the defect indicates the sign of the winding number, and opposite winding numbers are marked by shaded and open circles. **(c)** $|\nabla \Phi_{\mathbf{Q}_y}|$, the magnitude of the phase gradient. The red dashed curve marks a one-dimensional region of high phase gradient, to emphasize the correlation between panels. The amplitude is highly suppressed in regions with rapidly changing phase. Topological defects with opposite winding numbers densely populate the low amplitude regions.

shows that the DW amplitude is suppressed in regions with rapidly changing phase.

4.3.4 TOPOLOGICAL DEFECTS

In Figure 4.6b, topological defects where the phase winds by 2π around a point, are also apparent. Most of these point defects come in nearby pairs of opposite winding number within the low amplitude regions. The density of these closely spaced pairs varies with the size of Λ (*i.e.* the q bandwidth, Equation 2.10) used to extract the DW structure in $\Psi_{\mathbf{Q}_\alpha}(\mathbf{r})$. These closely spaced pairs likely do not represent significant features of the DW structure. However, in addition, there is a moderate concentration

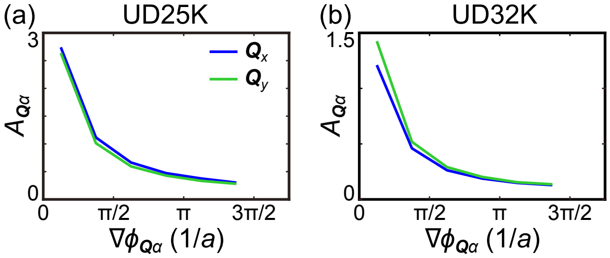

Figure 4.7: *d*FF DW correlation between amplitude and phase gradient. $A_{\mathbf{Q}_\alpha}$, the amplitude of the demodulated *d*FF DW $\Psi_{\mathbf{Q}_\alpha}$, averaged over regions with different phase gradients, $|\nabla\Phi_{\mathbf{Q}_\alpha}|$, shown for both x (blue) and y (green) direction modulations in UD25K (a) and UD32K (b). Regions with rapidly changing *d*FF DW phase have strongly suppressed amplitude.

of separated topological defects. In UD25K, we find eight (ten) topological defects that are not paired with a defect of opposite winding number within the same low-amplitude region for the x (y) direction *d*FF DW within the $85a \times 85a$ field of view, corresponding to approximately 3 defects in a 256 px crop. These observations are similar to reports in Bi2212 [123, 66].

4.4 Synthetic data generation

To generate a training data set resembling the experimental data, we use real data for the $\Delta(\mathbf{r})$ images and simulate corresponding synthetic $\mathcal{R}_s(\mathbf{r})$[**] images of x-direction *d*FF DW according to each hypothesis. To achieve a large and diverse data set, instances of $\Delta(\mathbf{r})$ overlap (producing more possible crops) and stochastic variations are built into the simulation. Conceptually the steps are (Figure 4.8):

[**] Here, the subscript s denotes synthetic data, not an *s*-form factor.

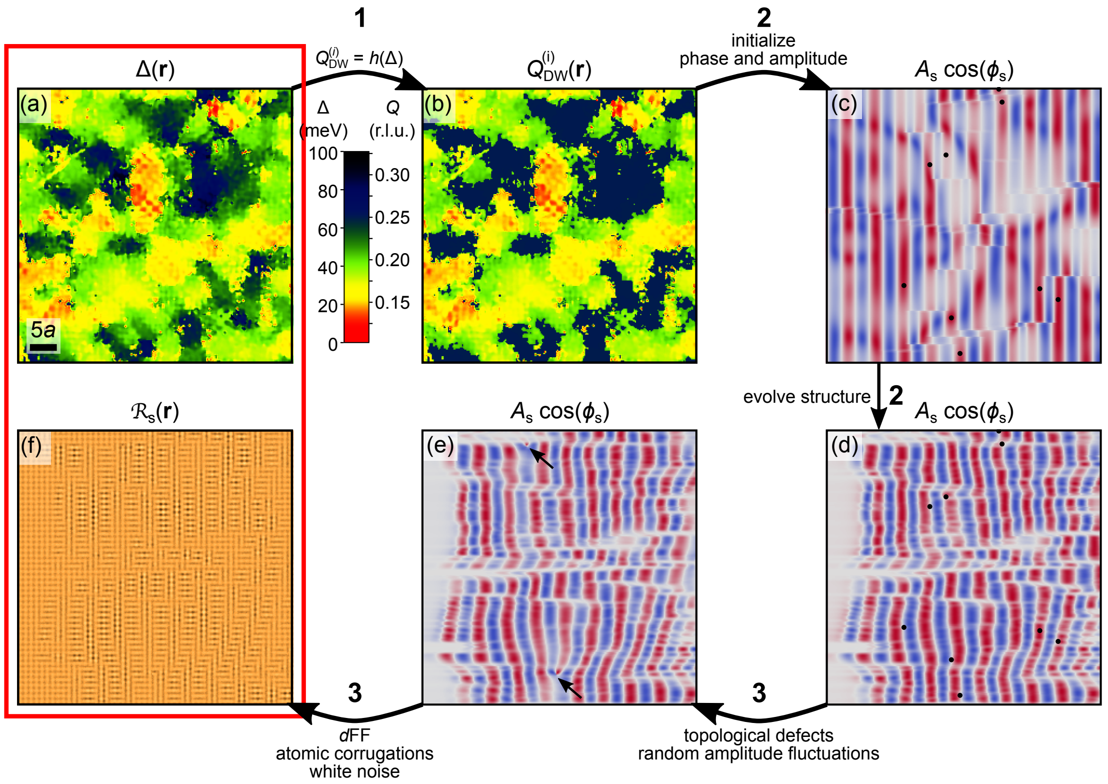

Figure 4.8: Overview of training data simulation. **(a)** An instance of $\Delta(\mathbf{r})$ for UD32K. **(b)** Ideal inhomogeneous $Q^{(i)}_{\mathrm{DW}}(\mathbf{r})$ generated using the piecewise hypothesis. **(c–e)** Simulation of the wave structure by initializing the amplitude and phase with the phase ϕ_s pinned at the sites marked by black dots (c), relaxing the structure toward $\partial_x \phi_s \approx Q^{(i)}_{\mathrm{DW}}$ (d), and adding in additional forms of disorder by hand (e). **(f)** Final synthetic data, including the *d*FF DW, atomic corrugations and white noise.

1. Use the hypothesis $Q_{\text{DW}}^{(i)} = h(\Delta)$ to calculate the ideal inhomogeneous wavevector map $Q_{\text{DW}}^{(i)}(\mathbf{r})$ from $\Delta(\mathbf{r})$ (Figure 4.8b).
2. Calculate the DW phase ϕ_s and amplitude A_s such that $Q_{\text{DW}}^{(i)} \approx \partial_x \phi_s$ in regions with significant A_s, where we are using the subscript s to indicate properties of the synthetic data (Figure 4.8c–d), not an s-form factor.
3. Add in the d-form factor, atomic corrugations, topological defects, random amplitude fluctuations and white noise (Figure 4.8e–f).

While the first step is straightforward, steps two and three involve key choices that govern how closely the synthetic $\mathcal{R}_s(\mathbf{r})$ data resembles the experimentally measured data. In the remainder of this section, we describe steps two and three in detail, and then evaluate the resulting synthetic data.

4.4.1 CALCULATING THE PHASE AND AMPLITUDE

The DW phase is defined by $Q_{\text{DW}}^{(i)} \approx \nabla \phi_s$, but this differential equation does not uniquely define ϕ_s and the exact solutions may not be physically realistic. Furthermore, we know from the real data that the spatial phase and amplitude are highly correlated. We therefore have developed a simulation that initializes the phase with disorder due to pinning (Figure 4.8c), then evolves the amplitude and phase in coordinated steps to improve fidelity of the wavevector, while maintaining correlated inhomogeneity.

To maintain correlation between $\nabla \phi_s$ and A_s, A_s is calculated from ϕ_s as:

$$A_s = A_s^{(\text{error})} \cdot A_s^{(\text{orient})} \tag{4.6}$$

The first term is a Lorentzian suppression of the amplitude based on the wavevector error $\partial_x \phi_s - Q_{\mathrm{DW}}^{(i)}$:

$$A_s^{(\mathrm{error})} = \frac{(c_1/2)^2}{(\partial_x \phi_s - Q_{\mathrm{DW}}^{(i)})^2 + (c_1/2)^2} \tag{4.7}$$

where $c_1 = 0.50\ 2\pi/a$. The second term $A_s^{(\mathrm{orient})}$ suppresses the amplitude if the orientation of the $\phi_s = 2\pi n$ wavefronts deviates from vertical:

$$A_s^{(\mathrm{orient})} = \frac{1}{2}\left[\cos\left(\pi\, \partial_y\, x_{\mathrm{crest}}\right) + 1.0\right], \tag{4.8}$$

where x_{crest} is the location of a wavefront marked by $\phi_s = 2\pi n$. Smoothing of the total A_s prevents rapid fluctuations.

We initialize ϕ_s using a constant wavevector and pinning the phase by the location of the nearest randomly chosen pinning site (Figure 4.9a–c). The number of pinning sites (black dots) is arbitrarily chosen between one and fifteen. This initialization by design generates one-dimensional regions of low A_s where ϕ_s jumps rapidly between domains. We iteratively improve ϕ_s by using three forces to evolve the positions of the wavefronts. The first force depends on the spacing between wavefronts to favor the ideal local wavevector. The second force acts like a stiffness, straightening out the wavefronts. The third force is a pinning force from the sites randomly chosen for the phase initialization. Crucially, the amplitude is included in the first two forces: (1)

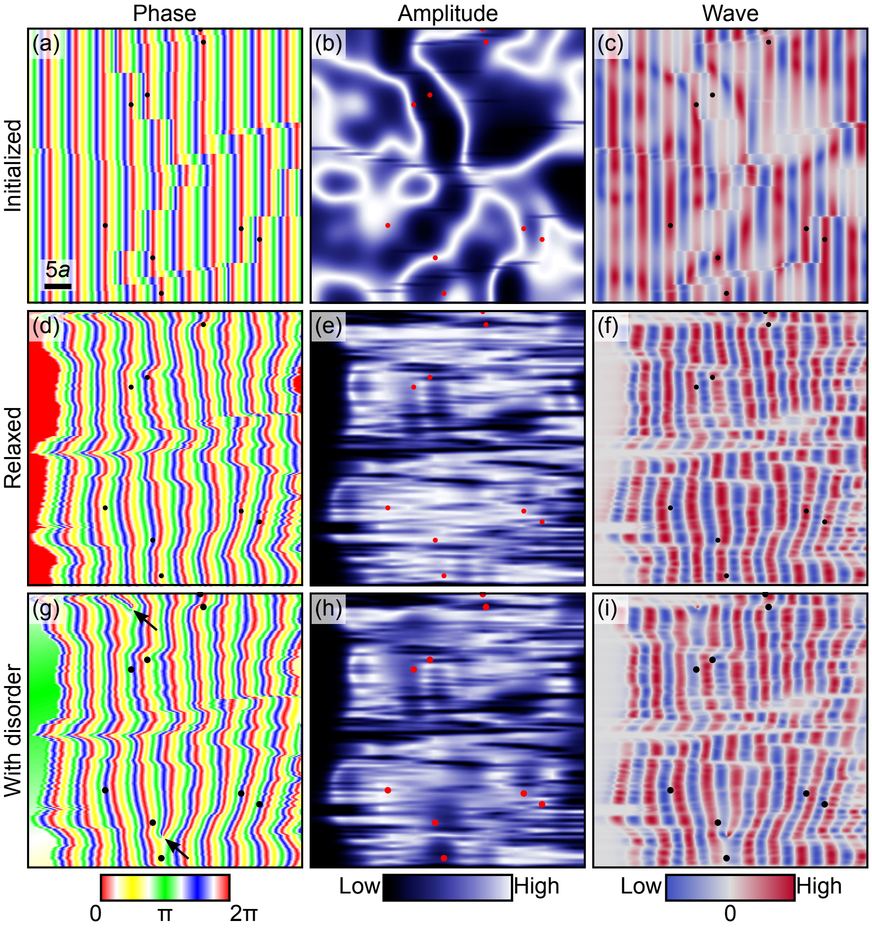

Figure 4.9: Stages of data simulation. Phase ϕ_s and amplitude A_s of the wave $A_s \cos[\phi_s]$ shown in Figure 4.8 at the three stages of simulation described in the text. The dots mark the locations of the pinning sites. The arrows mark the locations of two topological defects added in by hand.

The amplitude is used to weight the forces from the left and right neighboring wavefronts, allowing the wavevector approximation to improve in patches of higher amplitude, while sacrificing low amplitude regions in between. (2) The straightening force is scaled by A_s, allowing regions to form with rapidly changing phase and strongly suppressed amplitude, as observed in the experimental data. Figure 4.9d–f shows the resulting relaxed DW structure. The wavefronts have become curved to accommodate both spatial variations in $Q_{\mathrm{DW}}^{(i)}$ and the pinning sites. The structure maintains one-dimensional regions of low A_s and high $\nabla\phi_s$.

4.4.2 ADDITIONAL FORMS OF DISORDER

In the final stage, we add in the dFF, atomic corrugations and additional forms of disorder present in $\mathcal{R}(\mathbf{r})$:

$$\mathcal{R}_s(\mathbf{r}) = a_d(\mathbf{r})A_s(\mathbf{r})\cos[\phi_s(\mathbf{r})] + A_{\mathrm{lat}}f_{\mathrm{lat}}(\mathbf{r}) + A_{\mathrm{noise}}\eta(\mathbf{r}) \tag{4.9}$$

The first term is the dFF DW, the second the atomic corrugations and the third normally distributed white noise, where $\eta(\mathbf{r})$ has a standard deviation of one.

For the dFF DW, we introduce up to four topological defects with arbitrarily chosen polarity (*i.e.* winding number of ±1) into $\phi_s(\mathbf{r})$, and we add in the dFF $a_d(\mathbf{r})$, using only the first Fourier components (mimicking the momentum bandwidth of the STM data, as described in Section 2.4).

Table 4.1: Quantifying disorder in real and synthetic data. Ratio of the mean square of the *d*FF DW modulations (P_{DW}) to the mean square of the atomic corrugations (P_{atoms}), ratio of the mean square white noise P_{white} to P_{DW}, and the asymmetry between the Bragg peaks α for each of the four data sets. The right column gives the range of values which can occur in the synthetic data.

Quantity	UD25K	UD32K	OPT35K	OD15K	Synthetic data	
$\frac{P_{\mathrm{DW}}}{P_{\mathrm{atoms}}}$	6.0	7.0	0.12	8.0	[0.50	13]
$\frac{P_{\mathrm{white}}}{P_{\mathrm{DW}}}$	0.054	0.022	0.41	$4.8 \cdot 10^{-3}$	$[5.3 \cdot 10^{-4}$	0.36]
α	0.84	4.9	2.1	1.1	[0.50	2.0]

In the atomic corrugations, we include an asymmetry α between the x and y directions:

$$f_{\mathrm{lat}}(\mathbf{r}) = \frac{\alpha}{N} \cos\left[G(x - x0)\right] + \frac{1}{N} \cos\left[G(y - y0)\right], \tag{4.10}$$

where $G = 2\pi/a$, $N = \sqrt{1+\alpha^2}$ normalizes the amplitude, and (x_0, y_0) the position of the lattice is chosen randomly.

The amplitudes of these three components (A_{lat}, A_{noise}, and the average $A_s(\mathbf{r})$) are chosen randomly within limits set to mimic the real data, and the overall amplitude of the data is normalized to have a standard deviation of one. Table 4.1 compares the mean square amplitudes $P = A^2/2$ observed in the real data to the ranges covered by the synthetic data.

In summary, the amplitudes of the DW, Bragg peaks, and white noise, as well as the concentration of topological defects are chosen to produce a range of variation in

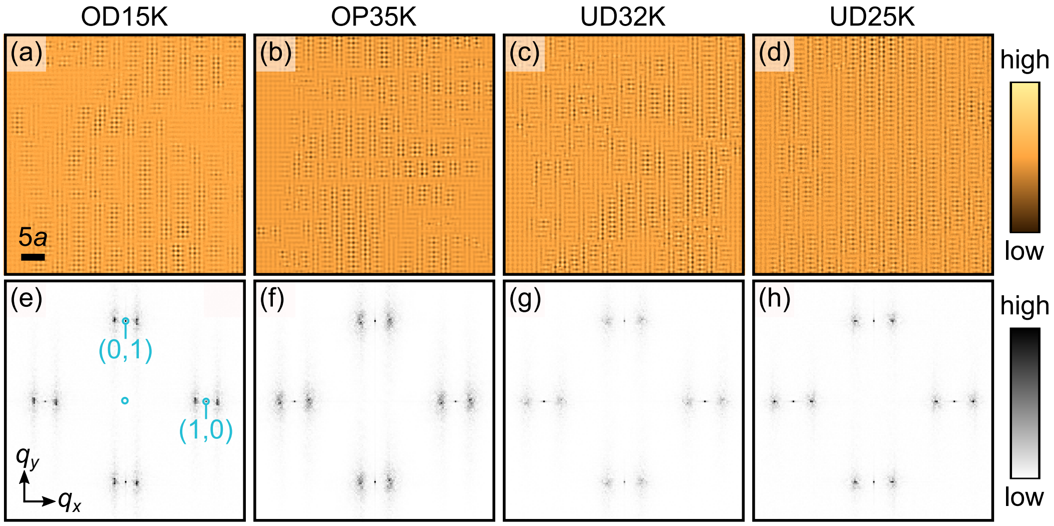

Figure 4.10: Instances of synthetic data $\mathcal{R}_s(\mathbf{r})$ (a–d) generated from the same $\Delta(\mathbf{r})$ maps shown in Figure 4.4 with the piecewise hypothesis. **(e–h)** Amplitude of the Fourier transform $\mathcal{R}_s(\mathbf{q})$ for (a–d).

the synthetic data with good representation of the regimes of the real data. We leave out three forms of disorder observed in the real data††. We do not include low q noise arising from localized defects, s- and s'FF modulations, and y-direction dFF modulations. Work by Zhang *et al.* [109] suggests that the CNNs may generalize to analyzing the x-direction modulations in the presence of superimposed y-direction modulations. The generalization to data with these excluded variations can and should be checked explicitly before classifying real data.

4.4.3 EVALUATING SYNTHETIC DATA

Figure 4.10 shows synthetic data for the same fields of view as the experimental data in Figure 4.4. The Fourier transforms $\mathcal{R}_s(\mathbf{q})$ have broad peaks resembling those in $\mathcal{R}(\mathbf{q})$. In real space, however, a direct visual comparison of $\mathcal{R}_s(\mathbf{r})$ to $\mathcal{R}(\mathbf{r})$ is misleading because we do not include y-direction modulations, s- and s'-form factor modulations, and low q noise. We therefore Fourier filter the peaks located near $(1 \pm \bar{Q}_{\text{DW}}, 0)$ and $(\pm\bar{Q}_{\text{DW}}, 1)$ r.l.u. in Figure 4.11 to isolate x-direction modulations of primarily dFF. The charge modulation patterns in the real and synthetic data show similar wavefront curvature and amplitude variations.

In addition to resembling the experimental data, the synthetic $\mathcal{R}_s$ should maintain good fidelity of $\partial_x \phi_s$ to $Q_{\text{DW}}^{(i)}$ in order for the relationship $Q_{\text{DW}}^{(i)} = h(\Delta)$ to be represented in the data. We test this fidelity by examining five instances of synthetic data simulated for the same field of view in Figure 4.12. The $\partial_x \phi_s$ maps show variations but all clearly maintain the spatial structure of $Q_{\text{DW}}^{(i)}$. More precisely, we can examine the distribution of $\partial_x \phi_s$ for each value of $Q_{\text{DW}}^{(i)}$ in a two-dimensional histogram, as shown in Figure 4.13. For a wave with the ideal inhomogeneous wave vector $\partial_x \phi_s = Q_{\text{DW}}^{(i)}$, all counts would fall along the diagonal dashed line. The synthetic data shows a spread of values, but maintains good correlation to $Q_{\text{DW}}^{(i)}$. We therefore

---

[††] We use the term disorder even though these components are in fact orders because they are signals which are not our signal of interest, and they have random variations.

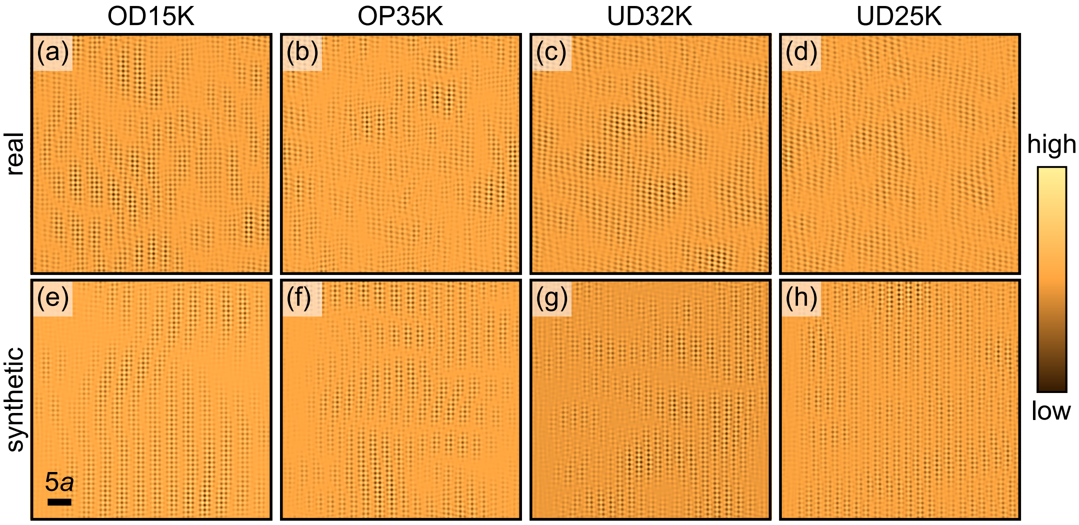

Figure 4.11: Comparison of modulations in real and synthetic data. Real ($\mathcal{R}(\mathbf{r})$ in a–d) and synthetic ($\mathcal{R}_s(\mathbf{r})$ in e–h) data from Figures 4.4 and 4.10 for identical fields of view Fourier filtered to isolate the x-direction charge modulations. Fourier filtering was done by masking each charge order peak located at $(1 \pm \bar{Q}_{\mathrm{DW}}, 0)$ and $(\pm\bar{Q}_{\mathrm{DW}}, 1)$ r.l.u., where $\bar{Q}_{\mathrm{DW}}$ is the sample average wave vector from Chapter 3, with a disk of radius $0.64\,\bar{Q}_{\mathrm{DW}}$ in the Fourier transformed data, then applying an inverse Fourier transform.

conclude that the $Q_{\mathrm{DW}}^{(i)} = h(\Delta)$ information remains embedded in the synthetic data.

4.4.4 SUMMARY

In summary, we include the following disorder in the synthetic training data:

- Inhomogeneous $Q_{\mathrm{DW}}(\Delta)$
- Phase disorder
- Amplitude variations (both correlated and uncorrelated to phase disorder)
- Topological defects
- White noise
- Phase of the lattice and asymmetry in the atomic corrugations

These forms of disorder include stochastic variations, enabling us to build a large and diverse training data set from the four $\Delta(\mathbf{r})$ maps measured via STM.

4.5 CNNS TO DISTINGUISH BETWEEN $Q_{\mathrm{DW}}(\Delta)$ HYPOTHESES

Five CNNs implemented in Keras [125] with a TensorFlow backend were trained with the following architecture:

1. Two dimensional convolutional layer with four filters and kernel size 4 by 4 px
2. Two dimensional convolutional layer with eight filters and kernel size 8 by 8 px
3. Two dimensional convolutional layer with 16 filters and kernel size 32 by 32 px
4. Fully connected layer with 16 neurons and sigmoid activation
5. Fully connected layer (output) with two neurons and softmax activation

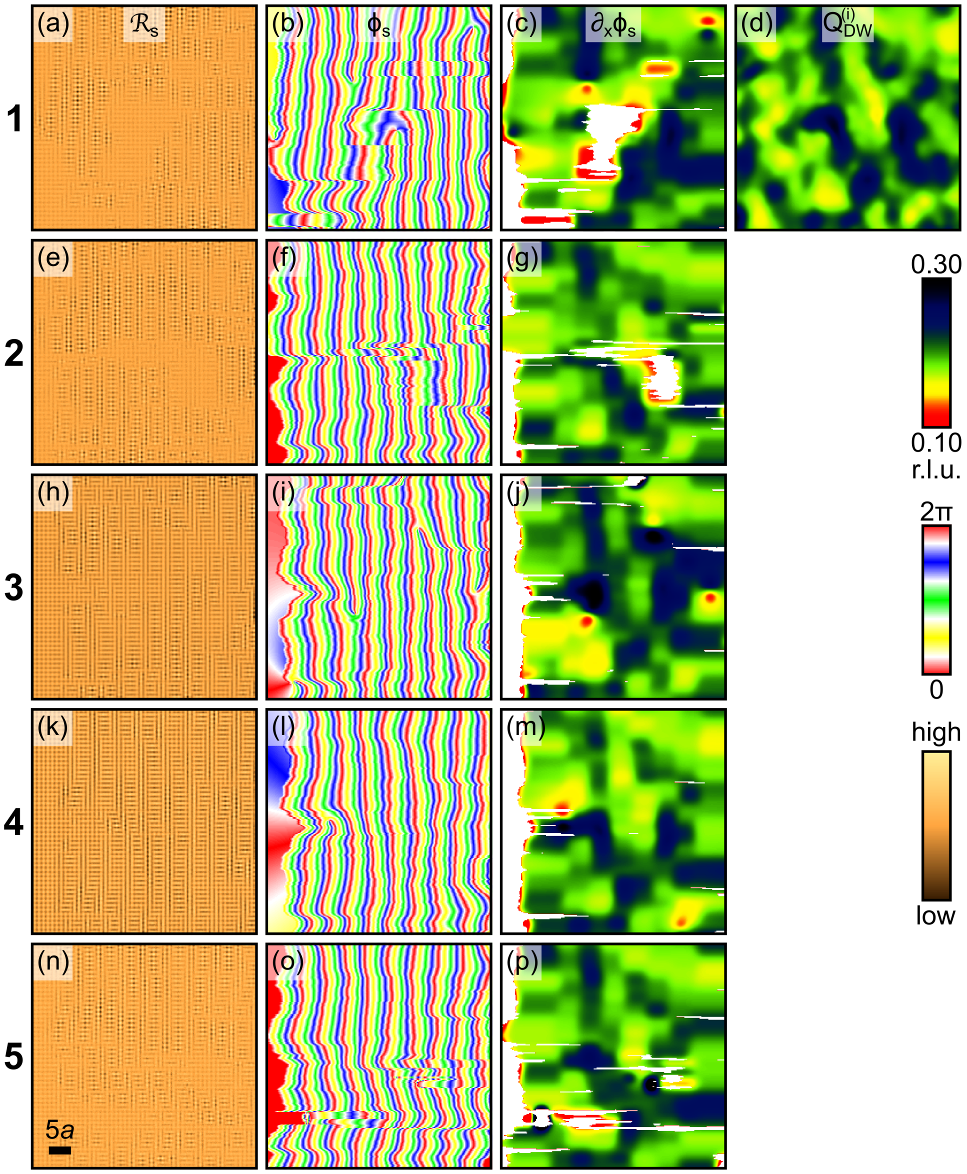

Figure 4.12: Multiple instances of synthetic data for the same field of view. $\mathcal{R}_s(\mathbf{r})$, $\phi_s(\mathbf{r})$ and $\partial_x\phi_s(\mathbf{r})$ from multiple simulations (rows 1-5) using the same $\Delta(\mathbf{r})$ instance from UD32K with the linear hypothesis. **(d)** Ideal inhomogeneous wave vector $Q_{\mathrm{DW}}^{(i)}(\mathbf{r})$ for comparison to $\partial_x\phi_s$. Both $Q_{\mathrm{DW}}^{(i)}(\mathbf{r})$ and $\partial_x\phi_s$ have been Gaussian smoothed with a filter width of a because we do not expect variations on smaller length scales to be relevant. White regions in the maps of $\partial_x\phi_s$ are regions with amplitude lower than 10% of the maximum amplitude.

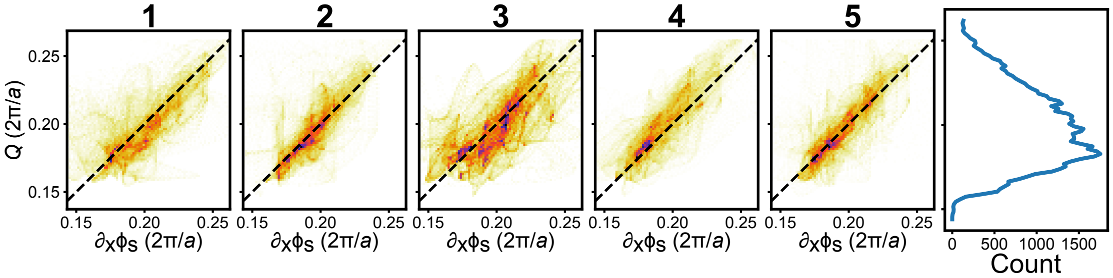

Figure 4.13: Q fidelity. **(a–e)** Two dimensional histograms of $Q_{\mathrm{DW}}^{(i)}(\mathbf{r})$ and $\partial_x \phi_s(\mathbf{r})$ for each of the instances of synthetic data shown in Figure 4.12. Some outliers occur outside of the range shown. **(f)** Histogram of $Q_{\mathrm{DW}}^{(i)}(\mathbf{r})$ values.

All convolutional layers use a hyperbolic tangent activation function and are followed by mean pooling with a pooling size of two by two pixels. The outputs of the final layer give the classification score for the two categories, respectively. The sum of the two outputs is constrained to 1.0 by the softmax activation such that on the linear output, 1.0 corresponds to confident classification as linear, whereas 0.0 corresponds to confident classification as piecewise and vice versa for the piecewise output.

We simulated 28000 instances of synthetic data (3500 per sample per hypothesis) for training and 4000 instances for validation. Training was done using categorical cross entropy for the loss function with L2 regularization included for all layers. Dropout at a rate of 0.25 was included just before the first fully connected layer to prevent overfitting. The classification accuracies during training, for the CNN that reached the highest accuracy (80.5% and 78.3% on the training and validation sets, respectively) are shown in Figure 4.14. With similar numbers of training epochs (1118, 1131, 1112 and 925), all five CNNs reached similar accuracies in the range of 77.0% to

80.5% on the training set and 73.3% to 78.3% on the validation set.

At face value, this performance seems poor, but we expect classification accuracy to be low on the samples with small Δ. In the regime $\Delta < \Delta_c$, the difference between the linear and piecewise hypotheses is small, 0.007 r.l.u. at $\Delta = 15$ meV (Figure 4.1). For comparison, the 256 px ($51.2a$) field of view of each instance gives a q-space resolution of approximately 0.02 r.l.u. We therefore expect that classification of OD15K- and OP35K-derived data should not be feasible, and classification of UD25K-derived data should be the most successful. Figure 4.14 breaks down the accuracies by sample and shows how the distribution of outputs varies from sample to sample. The CNNs perform well on UD25K-derived data with 98.8(8)% training accuracy and 98(1)% validation accuracy (the error represents the standard deviation of the five CNNs) and output values sharply peaked near maximum confidence for the correct hypothesis. The CNNs perform poorly on data from the three remaining samples with low accuracies and broadly distributed output values. Surprisingly, however, the performance is not significantly better on UD32K than on OD15K and OPT35K, and the accuracies on OD15K and OPT35K are significantly above random (0.5). The explanations behind these two observations remain to be determined.

Finally, we note that while the similarity between training and validation accuracies suggest little overfitting, because the validation and training data instances have overlapping fields of view, overfitting to specific features resulting from the spatial structure of the $\Delta(\mathbf{r})$ maps is not ruled out.

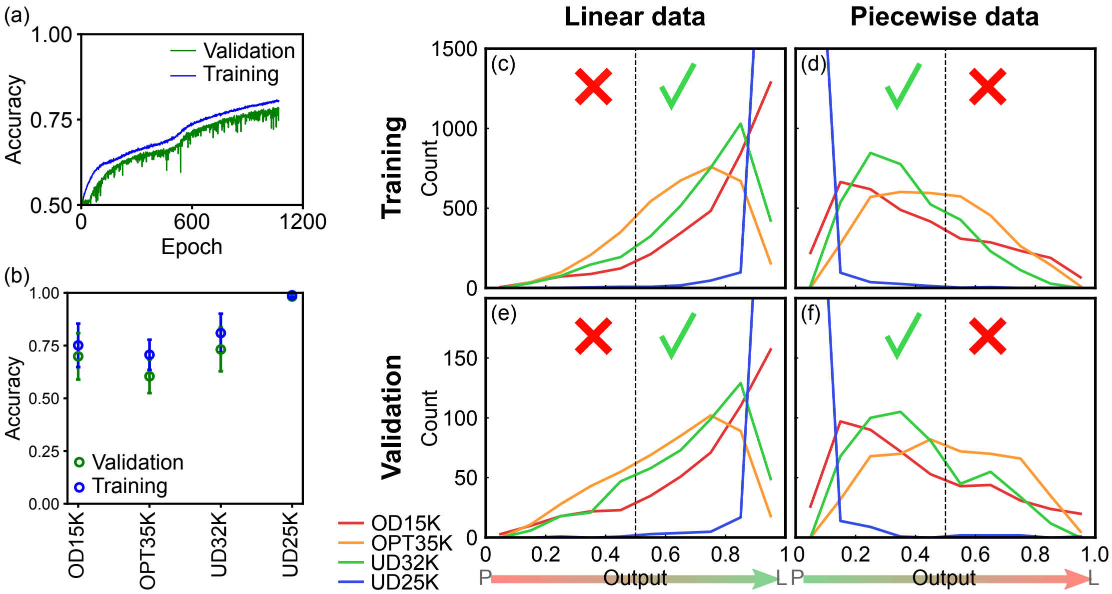

Figure 4.14: CNN performance. **(a)** Evolution of training and validation accuracies during the training process, shown for the CNN which reached the highest accuracy. **(b)** Training and validation accuracies for data from each sample. The points and error bars represent the average and standard deviation of the accuracies from all five CNNs. **(c–f)** Histograms of outputs from the CNN with the highest accuracy. 1.0 represents maximum confidence that the input represents the linear hypothesis, and 0.0 represents maximum confidence that the input represents the piecewise hypothesis. The colored curves are for input data derived from $\Delta(\mathbf{r})$ data on the four different samples. The four panels individually represent data simulated according to each hypothesis (columns) and as part of the training and validation sets (rows). In all panels the UD25K histogram is truncated in the bin with maximum counts.

4.6 DISCUSSION

While the CNNs' performances demonstrate that they are capable of discerning the two $Q_{\mathrm{DW}}(\Delta)$ hypotheses, we cannot yet conclude that the CNNs are classifying a local correlation between the dFF DW structure and the local gap. The hypothesis used to generate an instance of synthetic data governs global properties, such as the distribution of wavevectors and the average wavevector, in addition to the local variations. Therefore, it is possible that the CNNs are classifying the data based on global properties of $\mathcal{R}_s(\mathbf{r})$.

While the architecture we use is suitable for classifying synthetic data whose underlying structure is known to follow one of our two hypotheses, the architecture should not be used to classify real data. Because it is feasible that neither of our hypotheses describe the experimental data, a third category is needed for anything that does not fall into the two hypothesis categories, as implemented by Zhang *et al.* [109].

In addition, to transition from classifying synthetic data to real data, it must be demonstrated that the CNNs can generalize. This can be tested both by generating synthetic data with new types of disorder not seen in training, and by filtering real data to systematically remove and/or tune the disorder.

Finally, let us return to the ultimate scientific objective of this work. We would like to differentiate between two physically different scenarios: Is there a commensurate to incommensurate transition, or does a single incommensurate instability underly the full doping range? We have chosen two specific $h(\Delta)$ functions to represent the two scenarios, but of course those are not the only possible choices. The parameters chosen may be incorrect, or the true $Q_{\mathrm{DW}}(\Delta)$ may be curved rather than piecewise linear. Further testing would be necessary to characterize the sensitivity of classification to these choices, but we expect some degree of sensitivity because the training data for each category was simulated according to a single $h(\Delta)$ function. Section 4.7 makes a proposal for diversifying training data to reduce sensitivity to unimportant variations in the functional form of the hypothesis. We also note that schemes other than classification can be used to avoid this problem.

4.7 HYPOTHESIS DIVERSIFICATION

The goal of this work is to test two hypotheses for the $Q^{(i)}_{\mathrm{DW}}(\Delta)$ relationship underlying the spatially inhomogeneous charge modulations in Bi2201. In functional form:

$$\begin{aligned}
Q^{(i)}_{\mathrm{DW}} &= h_l(\Delta) &&= m_l\Delta + b_l && \text{linear hypothesis} \\
Q^{(i)}_{\mathrm{DW}} &= h_p(\Delta) &&= \begin{cases} 0.25 \text{ r.l.u.} & \text{for } \Delta \geq \delta_c \\ m_p\Delta + b_p & \text{for } \Delta < \delta_c \end{cases} && \text{piecewise hypothesis}
\end{aligned}$$

represent the hypotheses of a single linearly evolving incommensurate wavevector and a commensurate to incommensurate transition, respectively. If we choose a set of parameters for each ((m_l, b_l) and (m_p, b_p, δ_c)) to use for generating the training data, then we train the NN to distinguish between the two specific functional forms. However, our scientific interest is based on the qualitative difference between the hypotheses, not the details of their parametrization. In this section, we develop a framework for diversifying the hypotheses used to generate the training data with the goal of emphasizing the qualitative difference between the hypotheses and reducing the influence of small changes in parameters.

The strategy involves two main steps:

1. Generate probability distribution functions (PDFs) for the hypothesis parameterizations.
2. For each instance of training data, randomly sample the relevant parameters from the appropriate PDF.

In this section, we describe the implementation for our question of distinguishing a linear $Q_{\mathrm{DW}}^{(i)}(\Delta)$ from a piecewise $Q_{\mathrm{DW}}^{(i)}(\Delta)$.[*]

4.7.1 Probability distribution functions

In building our hypothesis PDFs, our Q_{DW} measurements from Chapter 3 provide a set of priors. Functions $h(\Delta)$ which do not describe the Q_{DW} data accurately are unlikely to be present in the experimental data and therefore do not need to be well-represented in our training data set. We therefore build our PDF based on fitting the Q_{DW} data and assuming normally distributed parameters. The PDFs for both

---

[*] In some sense, this strategy is no different from adding noise and other stochastic variations into training data, as is common. The key here is to extend the diversity to the scientific hypothesis being categorized.

hypotheses are displayed in Figure 4.15. The remainder of this subsection describes the technical details of building the PDFs.

Obtaining the linear hypothesis PDF $g_l(m_l, b_l)$ is straight forward. Linear regression on the Q_{DW} data produces mean values $\bar{m}_l$ and $\bar{b}_l$, as well as a covariance matrix Σ, which together define a normal distribution:

$$g_l(m_l, b_l) = \det\left(2\pi\Sigma\right)^{-1/2} \exp\left[-\frac{1}{2}\sum_{i,j}\Sigma_{ij}(p_i - \bar{p}_i)(p_j - \bar{p}_j)\right] \tag{4.11}$$

where $\mathbf{p} = (m_l, b_l)$ is a vector of the parameters, and $\bar{p} = (\bar{m}_l, \bar{b}_l)$. Figure 4.15 shows g_l and a set of parameters randomly sampled from the distribution.

Building a piecewise hypothesis PDF $g_p(m_p, b_p, \delta_c)$ is more complex because of the two-stage regression procedure described in Section 3.12. Based on these steps, we express g_p as a product of a conditional probability distribution and a prior:

$$g_p(m_p, b_p, \delta_c) = g_{\Delta}(\delta_c | m_p, b_p) g_{M,B}(m_p, b_p) \tag{4.12}$$

We obtain $g_{M,B}(m_p, b_p)$ by fitting the linear region of the data, as in Equation 4.11. For the conditional probability, we assume a normal distribution of δ_c for each choice of (m_p, b_p). We perform regressions to determine the mean and standard deviation for a grid of (m_p, b_p) values, and then interpolate to obtain a continuous $g_{\Delta}(\delta_c | m_p, b_p)$. Figure 4.15 shows the result. It is worth noting the strong cross-correlations among the parameters.

4.7.2 Distinguishability of the hypotheses

In Section 4.5, we considered the distinguishability of our two hypotheses. Figure 4.15b shows a set of diversified hypotheses sampled from the linear and piecewise PDFs. In this scheme, some of the linear hypotheses represented in the training data will over certain ranges of Δ be more likely to represent the piecewise rather than linear hypothesis. We therefore expect that the CNN should not be able to reach 100% training accuracy.

4.7.3 Contrast set

In Section 4.6, we discussed the need for a contrast category to represent a none-of-the-above classification. The hypothesis PDFs provide a natural framework for generating a contrast set. Once our hypothesis parameters are far enough from the optimal values, the functions they generate do not represent our physical hypotheses. Therefore, the contrast set can be populated by sampling from inverse PDFs over some domain. Alternatively, to reduce overlap between the contrast set and the hypothesis sets, one can use the PDF value to define a domain far from the optimal parametrization from which to sample parameters for the contrast set. In addition, one may be able to specifically exclude alternative hypotheses by generating PDFs for their parameterizations and including them in the contrast set. Overall, we expect this type

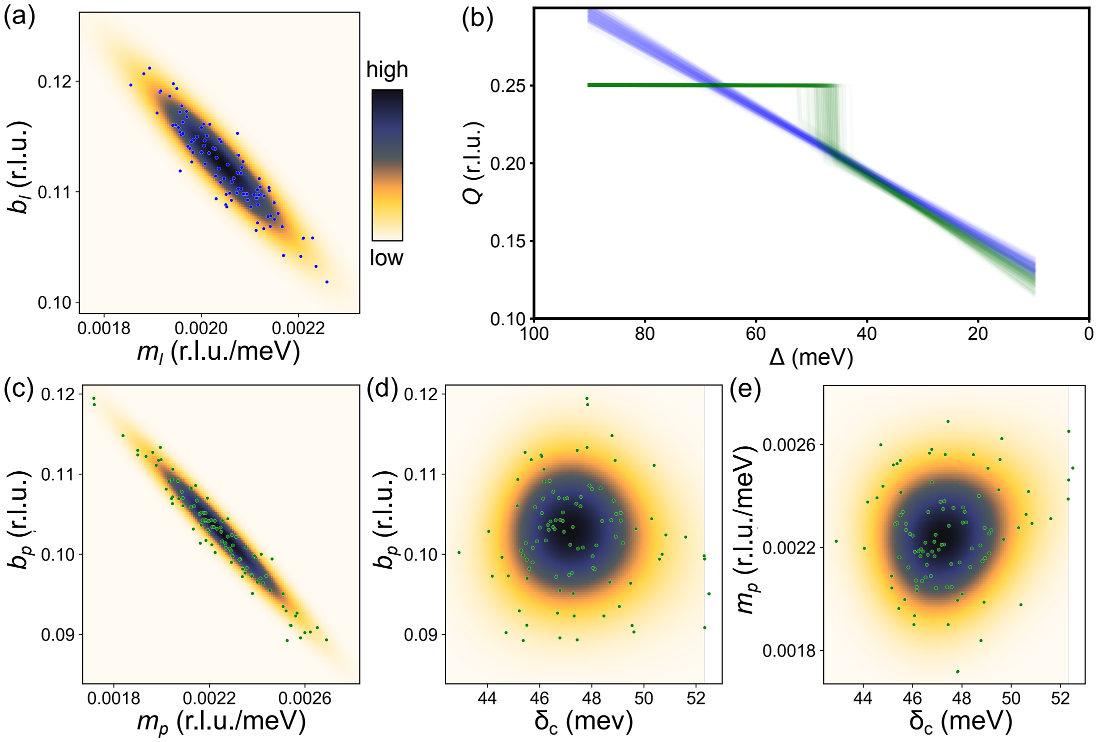

Figure 4.15: Stochastically sampling hypotheses. (a) The colorscale indicates the PDF for the linear hypothesis parameters. The points are randomly sampled parameter sets. **(b)** Randomly sampled linear (blue) and piecewise (green) hypotheses, corresponding to the data points shown in (a) and (c–e). **(c–e)** Projections of the PDF for the piecewise hypothesis parameters. The points are randomly sampled parameter sets.

of contrast set to help make the hypothesis classifications more precise.

4.8 CONCLUSIONS

In summary, we have quantified the inhomogeneous *d*FF DW structure in Bi2201, we have reported an algorithm to simulate realistic synthetic data, and we have demonstrated that convolutional neural networks are capable of distinguishing between different models underlying the local variations in the modulation wave vector. Multiple improvements are necessary before these types of CNNs can be used to analyze experimental data, and it remains to be determined if the classifications here are based on detecting local correlations. While we have emphasized the importance of this strategy for empowering discoveries related to the electronic structure of quantum materials, the potential applications are much broader as many fields rely on imaging experiments.

A

Python package for DW simulation

The bsco_dw package (located on the Hoffman lab's Google drive) contains tools for simulating the Bi2201 *d*FF DW. Most importantly, the simulate_dw module has the functions used for the simulations, and the hyp module has the tools used to randomly sample parameters for hypothesis diversification.

Parameter sampling for hypothesis diversification is done using the appropriate methods within the PDF_bivariate (for linear), PDF_multivariate (for multivariate hypotheses, with a covariance matrix describing the cross correlation between all parameters), or PDF_piecewise classes.

References

1. J. G. Bednorz and K. A. Müller, Zeitschrift für Phys. B Condens. Matter **64**, 189 (1986).
2. P. A. Lee, N. Nagaosa, and X.-G. Wen, Rev. Mod. Phys. **78**, 17 (2006).
3. J. Mannhart and D. G. Schlom, Science **327**, 1607 (2010).
4. H. Y. Hwang, Y. Iwasa, M. Kawasaki, B. Keimer, N. Nagaosa, and Y. Tokura, Nat. Mater. **11**, 103 (2012).
5. W. Witczak-Krempa, G. Chen, Y. B. Kim, and L. Balents, Annu. Rev. Condens. Matter Phys. **5**, 57 (2014).
6. B. Keimer, S. A. Kivelson, M. R. Norman, S. Uchida, and J. Zaanen, Nature **518**, 179 (2015).
7. E. Fradkin, in *Modern Theories of Many-Particle Systems in Condensed Matter Physics*, Lecture Notes in Physics, Vol. 843, edited by D. C. Cabra, A. Honecker, and P. Pujol (Springer Berlin Heidelberg, Berlin, Heidelberg, 2012) pp. 53–116.
8. E. Fradkin, S. A. Kivelson, and J. M. Tranquada, Rev. Mod. Phys. **87**, 457 (2015).
9. N. F. Mott, Proc. Phys. Soc. Sect. A **62**, 416 (1949).
10. G. Binning and H. Rohrer, Helv. Phys. Acta **55**, 726 (1982).
11. J. L. Tallon, J. W. Loram, G. V. M. Williams, J. R. Cooper, I. R. Fisher, J. D. Johnson, M. P. Staines, and C. Bernhard, Phys. status solidi **215**, 531 (1999).
12. B. Michon, C. Girod, S. Badoux, J. Kačmarčík, Q. Ma, M. Dragomir, H. A. Dabkowska, B. D. Gaulin, J.-S. Zhou, S. Pyon, T. Takayama, H. Takagi, S. Verret, N. Doiron-Leyraud, C. Marcenat, L. Taillefer, and T. Klein, Nature **567**, 218 (2019).

13. H. Murayama, Y. Sato, R. Kurihara, S. Kasahara, Y. Mizukami, Y. Kasahara, H. Uchiyama, A. Yamamoto, E.-G. Moon, J. Cai, J. Freyermuth, M. Greven, T. Shibauchi, and Y. Matsuda, ArXiv e-prints (2018), arXiv:1805.00276 .

14. Y. Sato, S. Kasahara, H. Murayama, Y. Kasahara, E.-G. Moon, T. Nishizaki, T. Loew, J. Porras, B. Keimer, T. Shibauchi, and Y. Matsuda, Nat. Phys. **13**, 1074 (2017).

15. A. Shekhter, B. J. Ramshaw, R. Liang, W. N. Hardy, D. A. Bonn, F. F. Balakirev, R. D. McDonald, J. B. Betts, S. C. Riggs, and A. Migliori, Nature **498**, 75 (2013).

16. I. M. Vishik, Rep. Prog. Phys. **81**, 062501 (2018).

17. L. Zhao, C. A. Belvin, R. Liang, D. A. Bonn, W. N. Hardy, N. P. Armitage, and D. Hsieh, Nat. Phys. **13**, 250 (2017).

18. R. Comin and A. Damascelli, Annu. Rev. Condens. Matter Phys. **7**, 369 (2016).

19. E. Berg, E. Fradkin, E.-A. Kim, S. A. Kivelson, V. Oganesyan, J. M. Tranquada, and S. C. Zhang, Phys. Rev. Lett. **99**, 127003 (2007).

20. E. Berg, E. Fradkin, S. A. Kivelson, and J. M. Tranquada, New J. Phys. **11**, 115004 (2009).

21. H.-D. Chen, O. Vafek, A. Yazdani, and S.-C. Zhang, Phys. Rev. Lett. **93**, 187002 (2004).

22. M. H. Hamidian, S. D. Edkins, S. H. Joo, A. Kostin, H. Eisaki, S. Uchida, M. J. Lawler, E.-A. Kim, A. P. Mackenzie, K. Fujita, J. Lee, and J. C. S. Davis, Nature **532**, 343 (2016).

23. S. D. Edkins, A. Kostin, K. Fujita, A. P. Mackenzie, H. Eisaki, S.-I. Uchida, S. Sachdev, M. J. Lawler, E.-A. Kim, J. C. S. Davis, and M. H. Hamidian, (2018), arXiv:1802.04673 .

24. E. W. Carlson, V. J. Emery, S. A. Kivelson, and D. Orgad, Concepts in High Temperature Superconductivity, in *Superconductivity*, edited by K. H. Bennemann and J. B. Ketterson (Springer, Berlin, Heidelberg, 2008) pp. 1225–1348.

25. S. Sachdev, E. Berg, S. Chatterjee, and Y. Schattner, Phys. Rev. B **94**, 115147 (2016).

26. L. Gao, Y. Y. Xue, F. Chen, Q. Xiong, R. L. Meng, D. Ramirez, C. W. Chu, J. H. Eggert, and H. K. Mao, Phys. Rev. B **50**, 4260 (1994).

27. A. Schilling, M. Cantoni, J. D. Guo, and H. R. Ott, Nature **363**, 56 (1993).

28. C. C. Tsuei and J. R. Kirtley, Rev. Mod. Phys. **72**, 969 (2000).

29. Y. Ando, Y. Hanaki, S. Ono, T. Murayama, K. Segawa, N. Miyamoto, and S. Komiya, Phys. Rev. B **61**, R14956 (2000).

30. R. G. Buckley, J. L. Tallon, I. W. M. Brown, M. R. Presland, N. E. Flower, P. W. Gilberd, M. Bowden, and N. B. Milestone, Physica C: Superconductivity **156**, 629 (1988).

31. J. L. Tallon, R. G. Buckley, P. W. Gilberd, and M. R. Presland, Physica C: Superconductivity **158**, 247 (1989).

32. A. Yamamoto, M. Onoda, E. Takayama-Muromachi, F. Izumi, T. Ishigaki, and H. Asano, Phys. Rev. B **42**, 4228 (1990).

33. W. W. Warren, R. E. Walstedt, G. F. Brennert, R. J. Cava, R. Tycko, R. F. Bell, and G. Dabbagh, Phys. Rev. Lett. **62**, 1193 (1989).

34. T. Timusk and B. Statt, Rep. Prog. Phys. **62**, 61 (1999).

35. M. R. Norman and C. Pépin, Rep. Prog. Phys. **66**, 1547 (2003).

36. S. Hüfner, M. A. Hossain, A. Damascelli, and G. A. Sawatzky, Rep. Prog. Phys. **71**, 062501 (2008).

37. K. Fujita, M. Hamidian, I. Firmo, S. Mukhopadhyay, C. K. Kim, H. Eisaki, S. Uchida, and J. C. Davis, Spectroscopic imaging STM: Atomic-scale visualization of electronic structure and symmetry in underdoped cuprates, in *Strongly Correlated Systems*, Springer Series in Solid-State Sciences, Vol. 180, edited by A. Avella and F. Mancini (Springer, Berlin, Heidelberg, 2015).

38. M. R. Norman, H. Ding, M. Randeria, J. C. Campuzano, T. Yokoya, T. Takeuchi, T. Takahashi, T. Mochiku, K. Kadowaki, P. Guptasarma, and D. G. Hinks, Nature **392**, 157 (1998).

39. K. M. Shen, F. Ronning, D. H. Lu, F. Baumberger, N. J. C. Ingle, W. S. Lee, W. Meevasana, Y. Kohsaka, M. Azuma, M. Takano, H. Takagi, and Z.-X. Shen, Science **307**, 901 (2005).

40. K. Tanaka, W. S. Lee, D. H. Lu, A. Fujimori, T. Fujii, Risdiana, I. Terasaki, D. J. Scalapino, T. P. Devereaux, Hussain, Z., and Z.-X. Shen, Science **314**, 1910 (2006).

41. A. Kanigel, M. R. Norman, M. Randeria, U. Chatterjee, S. Souma, A. Kaminski, H. M. Fretwell, S. Rosenkranz, M. Shi, T. Sato, T. Takahashi, Z. Z. Li, H. Raffy, K. Kadowaki, D. Hinks, L. Ozyuzer, and J. C. Campuzano, Nat. Phys. **2**, 447 (2006).

42. A. Kanigel, U. Chatterjee, M. Randeria, M. R. Norman, S. Souma, M. Shi, Z. Z. Li, H. Raffy, and J. C. Campuzano, Phys. Rev. Lett. **99**, 157001 (2007).

43. M. Platé, J. D. F. Mottershead, I. S. Elfimov, D. C. Peets, R. Liang, D. A. Bonn, W. N. Hardy, S. Chiuzbaian, M. Falub, M. Shi, L. Patthey, and A. Damascelli, Phys. Rev. Lett. **95**, 077001 (2005).

44. B. Vignolle, A. Carrington, R. A. Cooper, M. M. J. French, A. P. Mackenzie, C. Jaudet, D. Vignolles, C. Proust, and N. E. Hussey, Nature **455**, 952 (2008).

45. J. M. Luttinger, Phys. Rev. **119**, 1153 (1960).

46. M. Oshikawa, Phys. Rev. Lett. **84**, 3370 (2000).

47. Y. He, Y. Yin, M. Zech, A. Soumyanarayanan, M. M. Yee, T. Williams, M. C. Boyer, K. Chatterjee, W. D. Wise, I. Zeljkovic, T. Kondo, T. Takeuchi, H. Ikuta, P. Mistark, R. S. Markiewicz, A. Bansil, S. Sachdev, E. W. Hudson, and J. E. Hoffman, Science **344**, 608 (2014).

48. K. Fujita, C. K. Kim, I. Lee, J. Lee, M. H. Hamidian, I. A. Firmo, S. Mukhopadhyay, H. Eisaki, S. Uchida, M. J. Lawler, E.-A. Kim, and J. Davis, Science **344**, 612 (2014).

49. T. Kondo, R. Khasanov, T. Takeuchi, J. Schmalian, and A. Kaminski, Nature **457**, 296 (2009).

50. T. Kondo, Y. Hamaya, A. D. Palczewski, T. Takeuchi, J. S. Wen, Z. J. Xu, G. Gu, J. Schmalian, and A. Kaminski, Nat. Phys. **7**, 21 (2011).

51. G.-Q. Zheng, P. L. Kuhns, A. P. Reyes, B. Liang, and C. T. Lin, Phys. Rev. Lett. **94**, 047006 (2005).

52. J. M. Tranquada, B. J. Sternlieb, J. D. Axe, Y. Nakamura, and S. Uchida, Nature **375**, 561 (1995).

53. J. E. Hoffman, E. W. Hudson, K. M. Lang, V. Madhavan, H. Eisaki, S. Uchida, and J. C. Davis, Science , 466 (2002).

54. C. Howald, H. Eisaki, N. Kaneko, M. Greven, and A. Kapitulnik, Phys. Rev. B **67**, 014533 (2003).

55. K. Fujita, M. H. Hamidian, S. D. Edkins, C. K. Kim, Y. Kohsaka, M. Azuma, M. Takano, H. Takagi, H. Eisaki, S.-I. Uchida, A. Allais, M. J. Lawler, E.-A. Kim, S. Sachdev, and J. C. S. Davis, Proc. Natl. Acad. Sci. USA **111**, E3026 (2014).

56. Y. Kohsaka, C. Taylor, K. Fujita, A. Schmidt, C. Lupien, T. Hanaguri, M. Azuma, M. Takano, H. Eisaki, H. Takagi, S. Uchida, and J. C. Davis, Science **315**, 1380 (2007).

57. R. Comin, R. Sutarto, F. He, E. H. da Silva Neto, L. Chauviere, A. Fraño, R. Liang, W. N. Hardy, D. A. Bonn, Y. Yoshida, H. Eisaki, A. J. Achkar, D. G. Hawthorn, B. Keimer, G. A. Sawatzky, and A. Damascelli, Nat. Mater. **14**, 796 (2015).

58. A. J. Achkar, F. He, R. Sutarto, C. McMahon, M. Zwiebler, M. Hücker, G. D. Gu, R. Liang, D. A. Bonn, W. N. Hardy, J. Geck, and D. G. Hawthorn, Nat. Mater. **15**, 616 (2016).

59. S. A. Kivelson, E. Fradkin, and V. J. Emery, Nature **393**, 550 (1998).

60. M. A. Metlitski and S. Sachdev, Phys. Rev. B **82**, 075128 (2010).

61. S. Sachdev and R. La Placa, Phys. Rev. Lett. **111**, 027202 (2013).

62. Y. Wang and A. Chubukov, Phys. Rev. B **90**, 035149 (2014).

63. K. B. Efetov, H. Meier, and C. Pépin, Nat. Phys. **9**, 442 (2013).

64. W. D. Wise, K. Chatterjee, M. C. Boyer, T. Kondo, T. Takeuchi, H. Ikuta, Z. Xu, J. Wen, G. D. Gu, Y. Wang, and E. W. Hudson, Nat. Phys. **5**, 213 (2009).

65. R. Comin, A. Frano, M. M. Yee, Y. Yoshida, H. Eisaki, E. Schierle, E. Weschke, R. Sutarto, F. He, A. Soumyanarayanan, Y. He, M. Le Tacon, I. S. Elfimov, J. E. Hoffman, G. A. Sawatzky, B. Keimer, and A. Damascelli, Science **343**, 390 (2014).

66. A. Mesaros, K. Fujita, S. D. Edkins, M. H. Hamidian, H. Eisaki, S.-I. Uchida, J. C. S. Davis, M. J. Lawler, and E.-A. Kim, Proc. Natl. Acad. Sci. USA **113**, 12661 (2016).

67. M. H. Hamidian, S. D. Edkins, C. K. Kim, J. C. Davis, A. P. Mackenzie, H. Eisaki, S. Uchida, M. J. Lawler, E.-A. Kim, S. Sachdev, and K. Fujita, Nat. Phys. **12**, 150 (2016).

68. S. Badoux, W. Tabis, F. Laliberté, G. Grissonnanche, B. Vignolle, D. Vignolles, J. Béard, D. A. Bonn, W. N. Hardy, R. Liang, N. Doiron-Leyraud, L. Taillefer, and C. Proust, Nature **531**, 210 (2016).

69. S. Badoux, S. A. A. Afshar, B. Michon, A. Ouellet, S. Fortier, D. LeBoeuf, T. P. Croft, C. Lester, S. M. Hayden, H. Takagi, K. Yamada, D. Graf, N. Doiron-Leyraud, and L. Taillefer, Phys. Rev. X **6**, 021004 (2016).

70. I. Zeljkovic, Z. Xu, J. Wen, G. Gu, R. S. Markiewicz, and J. E. Hoffman, Science **337**, 320 (2012).

71. K. McElroy, D.-H. Lee, J. E. Hoffman, K. M. Lang, J. Lee, E. W. Hudson, H. Eisaki, S. Uchida, and J. C. Davis, Phys. Rev. Lett. **94**, 197005 (2005).

72. Y. Fei, K. Bu, W. Zhang, Y. Zheng, X. Sun, Y. Ding, X. Zhou, and Y. Yin, Sci. China Phys. Mech. Astron. **61**, 127404 (2018).

73. K. McElroy, J. Lee, J. A. Slezak, D.-H. Lee, H. Eisaki, S. Uchida, and J. C. Davis, Science **309**, 1048 (2005).

74. A. Piriou, N. Jenkins, C. Berthod, I. Maggio-Aprile, and Ø. Fischer, Nat. Commun. **2**, 221 (2011).

75. G. Kinoda, H. Mashima, K. Shimizu, J. Shimoyama, K. Kishio, and T. Hasegawa, Phys. Rev. B **71**, 020502(R) (2005).

76. C. J. Chen, *Introduction to Scanning Tunneling Microscopy* (Oxford University Press USA - OSO, Cary, 1993).

77. M. J. Lawler, K. Fujita, J. Lee, A. R. Schmidt, Y. Kohsaka, C. K. Kim, H. Eisaki, S. Uchida, J. C. Davis, J. P. Sethna, and E.-A. Kim, Nature **466**, 347 (2010).

78. K. Fujita, I. Grigorenko, J. Lee, M. Wang, J. X. Zhu, J. C. Davis, H. Eisaki, S. Uchida, and A. V. Balatsky, Phys. Rev. B **78**, 054510 (2008).

79. T. Hanaguri, Y. Kohsaka, J. C. Davis, C. Lupien, I. Yamada, M. Azuma, M. Takano, K. Ohishi, M. Ono, and H. Takagi, Nat. Phys. **3**, 865 (2007).

80. M. A. Metlitski and S. Sachdev, New J. Phys. **12**, 105007 (2010).

81. D. Podolsky, E. Demler, K. Damle, and B. I. Halperin, Phys. Rev. B **67**, 094514 (2003).

82. G. Ghiringhelli, M. Le Tacon, M. Minola, S. Blanco-Canosa, C. Mazzoli, N. B. Brookes, G. M. De Luca, A. Frano, D. G. Hawthorn, F. He, T. Loew, M. M. Sala, D. C. Peets, M. Salluzzo, E. Schierle, R. Sutarto, G. A. Sawatzky, E. Weschke, B. Keimer, and L. Braicovich, Science **337**, 821 (2012).

83. E. H. da Silva Neto, P. Aynajian, A. Frano, R. Comin, E. Schierle, E. Weschke, A. Gyenis, J. Wen, J. Schneeloch, Z. Xu, S. Ono, G. Gu, M. Le Tacon, and A. Yazdani, Science **343**, 393 (2014).

84. W. L. McMillan, Phys. Rev. B **14**, 1496 (1976).

85. B. Vignolle, A. Carrington, R. A. Cooper, M. M. J. French, A. P. Mackenzie, C. Jaudet, D. Vignolles, C. Proust, and N. E. Hussey, Nature **455**, 952 (2008).

86. C. Collignon, S. Badoux, S. A. A. Afshar, B. Michon, F. Laliberté, O. Cyr-Choinière, J.-S. Zhou, S. Licciardello, S. Wiedmann, N. Doiron-Leyraud, and L. Taillefer, Phys. Rev. B **95**, 224517 (2017).

87. A. Legros, S. Benhabib, W. Tabis, F. Laliberté, M. Dion, M. Lizaire, B. Vignolle, D. Vignolles, H. Raffy, Z. Z. Li, P. Auban-Senzier, N. Doiron-Leyraud, P. Fournier, D. Colson, L. Taillefer, and C. Proust, Nat. Phys. **15**, 142 (2019).

88. R. A. Cooper, Y. Wang, B. Vignolle, O. J. Lipscombe, S. M. Hayden, Y. Tanabe, T. Adachi, Y. Koike, M. Nohara, H. Takagi, C. Proust, and N. E. Hussey, Science **323**, 603 (2009).

89. I. Božović, X. He, J. Wu, and A. T. Bollinger, Nature **536**, 309 (2016).

90. J. Wu, A. T. Bollinger, X. He, and I. Božović, Nature **547**, 432 (2017).

91. F. Mahmood, X. He, I. Božović, and N. P. Armitage, Phys. Rev. Lett. **122**, 027003 (2019).

92. N. R. Lee-Hone, J. S. Dodge, and D. M. Broun, Phys. Rev. B **96**, 024501 (2017).

93. N. R. Lee-Hone, V. Mishra, D. M. Broun, and P. J. Hirschfeld, Phys. Rev. B **98**, 054506 (2018).

94. M. Le Tacon, M. Minola, D. C. Peets, M. Moretti Sala, S. Blanco-Canosa, V. Hinkov, R. Liang, D. A. Bonn, W. N. Hardy, C. T. Lin, T. Schmitt, L. Braicovich, G. Ghiringhelli, and B. Keimer, Phys. Rev. B **88**, 020501(R) (2013).

95. M. P. M. Dean, G. Dellea, R. S. Springell, F. Yakhou-Harris, K. Kummer, N. B. Brookes, X. Liu, Y. J. Sun, J. Strle, T. Schmitt, L. Braicovich, G. Ghiringhelli, I. Božović, and J. P. Hill, Nat. Mater. **12**, 1019 (2013).

96. D. F. Mross and T. Senthil, Phys. Rev. Lett. **108**, 267001 (2012).

97. D. F. Mross and T. Senthil, Phys. Rev. B **86**, 115138 (2012).

98. Y. Y. Peng, R. Fumagalli, Y. Ding, M. Minola, S. Caprara, D. Betto, M. Bluschke, G. M. De Luca, K. Kummer, E. Lefrançois, M. Salluzzo, H. Suzuki,

M. Le Tacon, X. J. Zhou, N. B. Brookes, B. Keimer, L. Braicovich, M. Grilli, and G. Ghiringhelli, Nat. Mater. **17**, 697 (2018).

99. Y. Y. Peng, M. Salluzzo, X. Sun, A. Ponti, D. Betto, A. M. Ferretti, F. Fumagalli, K. Kummer, M. Le Tacon, X. J. Zhou, N. B. Brookes, L. Braicovich, and G. Ghiringhelli, Phys. Rev. B **94**, 184511 (2016).

100. W. D. Wise, M. C. Boyer, K. Chatterjee, T. Kondo, T. Takeuchi, H. Ikuta, Y. Wang, and E. W. Hudson, Nat. Phys. **4**, 696 (2008).

101. P. Cai, W. Ruan, Y. Peng, C. Ye, X. Li, Z. Hao, X. Zhou, D.-H. Lee, and Y. Wang, Nat. Phys. **12**, 1047 (2016).

102. N. Miyakawa, P. Guptasarma, J. F. Zasadzinski, D. G. Hinks, and K. E. Gray, Phys. Rev. Lett. **80**, 157 (1998).

103. P. J. White, Z.-X. Shen, C. Kim, J. M. Harris, A. G. Loeser, P. Fournier, and A. Kapitulnik, Phys. Rev. B **54**, R15669(R) (1996).

104. J. M. Harris, Z.-X. Shen, P. J. White, D. S. Marshall, M. C. Schabel, J. N. Eckstein, and I. Bozovic, Phys. Rev. B **54**, R15665(R) (1996).

105. H. Ding, J. R. Engelbrecht, Z. Wang, J. C. Campuzano, S.-C. Wang, H.-B. Yang, R. Rogan, T. Takahashi, K. Kadowaki, and D. G. Hinks, Phys. Rev. Lett. **87**, 227001 (2001).

106. Q.-H. Wang and D.-H. Lee, Phys. Rev. B **67**, 020511(R) (2003).

107. J. E. Hoffman, K. McElroy, D.-H. Lee, K. M. Lang, H. Eisaki, S. Uchida, and J. C. Davis, Science **297**, 1148 (2002).

108. K. McElroy, R. W. Simmonds, J. E. Hoffman, D.-H. Lee, J. Orenstein, H. Eisaki, S. Uchida, and J. C. Davis, Nature **422**, 592 (2003).

109. Y. Zhang, A. Mesaros, K. Fujita, S. D. Edkins, M. H. Hamidian, K. Ch'ng, H. Eisaki, S. Uchida, J. C. S. Davis, E. Khatami, and E.-A. Kim, ArXiv e-prints (2018), arXiv:1808.00479 .

110. H. Zhao, Z. Ren, B. Rachmilowitz, J. Schneeloch, R. Zhong, G. Gu, Z. Wang, and I. Zeljkovic, Nat. Mater. **18**, 103 (2019).

111. D. Chowdhury and S. Sachdev, Phys. Rev. B **90**, 245136 (2014).

112. P. Bak, Rep. Prog. Phys. **45**, 587 (1982).

113. J. W. Alldredge, J. Lee, K. McElroy, M. Wang, K. Fujita, Y. Kohsaka, C. Taylor, H. Eisaki, S. Uchida, P. J. Hirschfeld, and J. C. Davis, Nat. Phys. **4**, 319 (2008).

114. C.-Z. Chang, J. Zhang, X. Feng, J. Shen, Z. Zhang, M. Guo, K. Li, Y. Ou, P. Wei, L.-L. Wang, Z.-Q. Ji, Y. Feng, S. Ji, X. Chen, J. Jia, X. Dai, Z. Fang, S.-C. Zhang, K. He, Y. Wang, L. Lu, X.-C. Ma, and Q.-K. Xue, Science **340**, 167 (2013).

115. P. Zhang, K. Yaji, T. Hashimoto, Y. Ota, T. Kondo, K. Okazaki, Z. Wang, J. Wen, G. D. Gu, H. Ding, and S. Shin, Science **360**, 182 (2018).

116. A. G. Kusne, T. Gao, A. Mehta, L. Ke, M. C. Nguyen, K.-M. Ho, V. Antropov, C.-Z. Wang, M. J. Kramer, C. Long, and I. Takeuchi, Sci. Rep. **4**, 6367 (2014).

117. M. Ziatdinov, O. Dyck, A. Maksov, X. Li, X. Sang, K. Xiao, R. R. Unocic, R. Vasudevan, S. Jesse, and S. V. Kalinin, ACS Nano **11**, 12742 (2017).

118. A. Bohrdt, C. S. Chiu, G. Ji, M. Xu, D. Greif, M. Greiner, E. Demler, F. Grusdt, and M. Knap, ArXiv e-prints (2018), arXiv:1811.12425 .

119. M. A. Nielsen, *Neural networkds and deep learning* (Determination Press, 2015) `http://www.neuralnetworksanddeeplearning.com`.

120. I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, 2016) `http://www.deeplearningbook.org`.

121. G. Cybenko, Math. Control. Signals, Syst. **2**, 303 (1989).

122. K. Hornik, M. Stinchcombe, and H. White, Neural Networks **2**, 359 (1989).

123. A. Mesaros, K. Fujita, H. Eisaki, S. Uchida, J. C. Davis, S. Sachdev, J. Zaanen, M. J. Lawler, and E.-A. Kim, Science **333**, 426 (2011).

124. R. Comin, R. Sutarto, E. H. da Silva Neto, L. Chauviere, R. Liang, W. N. Hardy, D. A. Bonn, F. He, G. A. Sawatzky, and A. Damascelli, Science **347**, 1335 (2015).

125. F. Chollet and *et al.*, Keras, `https://keras.io` (2015)