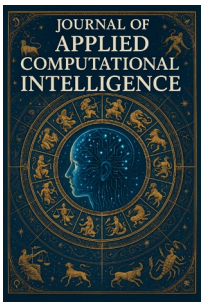
UNIVERSITY OF WEST ATTICA

eISSN: 3136-6937

Journal of Applied Computational Intelligence

Vol. 4, No. 1, January, 2026



Delay-Doppler Sparsity-Constrained U-Net for OTFS Channel Estimation under Fractional Doppler

Riccardo Amato[1,*], Stefano Di Marco[1], Antonio Ferrara[2], Emiliano Gentile[3]

[1] *Faculty of Information Engineering, Politecnico di Torino, Turin, 10129, Italy*
[2] *Faculty of Computer Science, University of Pisa, Pisa, 56126, Italy*
[3] *Faculty of Information Engineering, University of Calabria, Rende, 87036, Italy*

Received 23 November 2025; Revised 19 Feburary 2026; Accepted 27 Feburary 2026; Published 01 March 2026

Abstract

Fractional Doppler distributes the response of a propagation path across neighboring delay-Doppler cells, complicating sparse channel estimation in orthogonal time-frequency space systems. This study presents a residual U-Net that combines local Doppler neighborhoods with observation consistency to reconstruct an effective sampled channel. Its four input channels contain the real and imaginary components of a pilot-based initialization, a visibility mask, and a noise scale map. A piecewise differentiable shrinkage stage uses shared neighborhood energy to regulate the reconstructed support while preserving the phase of retained coefficients. Five estimator configurations are compared within a common evaluation specification. At a nominal signal-to-noise ratio of 20 dB, the reported normalized mean square error is −22.4 dB for the constrained configuration and −18.8 dB for the unconstrained U-Net, a separation of 3.6 dB. The corresponding support precision and recall are 0.91 and 0.89, compared with 0.82 and 0.86 for the unconstrained configuration. Ablation, offset, and path-density comparisons examine the relationship between neighborhood sharing and leakage reconstruction. The contribution concerns the integration of residual reconstruction, noise-conditioned neighborhood shrinkage, and an explicit observation penalty under integer-aligned delays. The reported coordinates describe the numerical comparison; the specified training allocation and device-level receiver validation remain separate implementation requirements. Practical assessment should examine pulse conventions, synchronization mismatch, fractional delay, and fixed-detector performance. These conditions define the scope of the present channel-domain interpretation.

Keywords: Fractional Doppler, OTFS Channel Estimation, Delay Doppler Sparsity, U Net, Observation Consistency

1. Introduction

High-mobility propagation challenges multicarrier communication through Doppler-induced interference and the need for accurate receiver channel estimation [1]. In orthogonal time-frequency space (OTFS) systems, time-frequency-domain pilot schemes provide an alternative observation representation for estimating channels with fractional delays and Dopplers [2]. Pulse shaping in the delay-Doppler domain, including designs based on Hermite basis functions for Zak-OTFS, provides a means of controlling the effective channel representation [3]. Spectrally efficient reference-signal designs combine pilot allocation with iterative channel estimation to address the acquisition overhead of OTFS modulation [4]. Fractional Doppler spreading is observed to extend across adjacent Doppler cells even in the presence of few strong scatterers in the physical environment. Therefore, channel estimation needs to find a suitable response neighborhood and reconstruct the complex coefficients in this neighborhood. Suppressing weak neighboring samples can eliminate the beneficial leakage energy, and accepting all neighboring samples will retain interference and noise.

* Corresponding author: amato.ric@polito.it

http://dx.doi.org/10.68406/jaci.2026.vol4iss1no8:96-110

➢ This is an open access article under the CC BY-NC-SA 4.0 license (https://creativecommons.org/licenses/by/4.0/).

© Authors retain all copyrights.

This difficulty becomes more pronounced at low pilot power, within narrow observation windows, or when weak path responses approach the local noise floor. Reliable reconstruction requires joint evaluation of neighborhood coherence and observation reliability. The largest coefficient is used as the basis for the above analysis, and the amplitude and phase pattern of the surrounding components show off-grid displacements. Therefore, a reconstruction model has been proposed that integrates local details, a larger Doppler context, and an explicit connection to the obtained pilot observations.

Sparse OTFS channel estimators exploit concentrated delay-Doppler responses, with sparsity-agnostic smoothed-zero-norm recovery providing one approach to coefficient reconstruction [5]. The accuracy depends on how closely the grid represents the physical propagation parameters. Fractional delay and Doppler estimation frameworks explicitly address the mismatch between continuous propagation parameters and the sampled OTFS representation [6]. If these samples are considered unrelated sparse entries, there will be a mismatch between the regularizer and the channel. A single-tone-based OTFS estimator recovers delay taps, integer and fractional Doppler components, and channel gains through sequential parameter estimation [7]. Multiscale attention residual networks have been investigated for channel computation, providing a learning-based approach to combining features at different scales [8]. The representation remains a central modeling choice in a neural estimator. A network trained on a small number of offsets may learn a convenient but incomplete leakage pattern. Such an estimate can reduce the channel tails of a detector. A single path can account for most of the channel energy and thus reduces the impact of a global error function on weak component attenuation. Therefore, in addition to reconstruction quality, neighborhood support, residual structure and sensitivity to out-of-distribution conditions should also be evaluated.

Attention residual U-Net architectures have been applied to channel estimation in massive multiple-input multiple-output (MIMO) free-space optical systems, providing a related example of encoder-decoder reconstruction [9]. The optical and unmanned-aerial-vehicle examples supply architecture-level context; the cited OTFS studies supply the directly relevant channel-estimation background. This selected coverage supports a scoped methodological comparison, with complex representations and boundary treatment requiring explicit design choices. Model-driven channel estimation networks have also been developed specifically for OTFS systems, connecting learned estimation with the underlying recovery model [10]. Bayesian recovery with Kronecker-structured sparsity has been studied for intelligent reflecting surface assisted multiple-input multiple-output channel estimation, illustrating a structured alternative to unstructured sparse modeling [11]. Both approaches depend on their modeling assumptions: operator errors produce an inaccurate acquisition model, while oversized neighborhoods can merge distinct path responses. Designed training data have been investigated as a means of improving the robustness of channel estimation neural networks [12]. A further difficulty is methodological transparency. Reliable performance assessment requires explicit baseline definitions, matched observation budgets, reproducible data partitions, and quantified uncertainty. Based on the above requirements, an estimator and an evaluation specification have been developed. The evaluation links reconstruction quality to support behavior, observation compatibility, and computational requirements.

To reconstruct the coherent leakage around a fractional Doppler path, a model is needed that combines spatial context with the reliability of pilot observations; thus, the proposed residual U-Net introduces a delay-Doppler neighborhood constraint and an observation consistency term to meet this requirement. The general design principle is to promote a small number of coherent response neighborhoods and keep the complex leakage coefficients within each selected neighborhood. The four input channels contain the real and imaginary initialization, the visibility mask, and the noise scale map; the latter two describe acquisition coverage and observation reliability. The architecture in this paper is an explicit observation model, and numerical analysis of the relationship between leakage geometry, support recovery, and computational requirements has been conducted. The remainder of the paper has the following structure. Section 2 introduces embedded pilot estimation, off-grid sparse recovery and learned reconstruction, and analyzes their assumptions and deficiencies. Section 3 defines the observation model, residual network, neighborhood constraints, and training objective. Section 4 introduces the evaluation plan and analyses the reconstruction error, support sensitivity, robustness and computational requirements. Section 5 presents the scope and limitations of the contribution, and the research required to verify the proposed mechanism in actual acquisition and deployment scenarios.

2. Related Work

2.1. Embedded Pilots and Fractional Doppler Estimation

Low-overhead OTFS transmission schemes have investigated channel estimation in the frequency or time domain as alternatives to delay-Doppler-domain acquisition [13]. For embedded-pilot acquisition, the protected region is specified in relation to the assumed effective channel spread. Integer Doppler reasoning may underestimate the range of interference in the path between Doppler grid points. For multiuser MIMO-OTFS with fractional delay and Doppler, cyclic-shift embedded pilots enable pilot reuse through the cyclic-shift orthogonality of Zadoff-Chu sequences [14]. The corresponding engineering trade-off is to reduce pilot energy and displaced data symbols and residual interference from the surrounding payload. Estimator comparisons require fixed acquisition resources or explicit accounting of resource

differences. Changes in guard region size alter observation quality and introduce an acquisition effect into the comparison of reconstruction algorithms. Windowed OTFS channel estimation studies show that appropriate time-frequency windows can suppress data-to-pilot interference in fractional Doppler channels [15]. Such processing changes the effective reconstruction target. A change in the receiver window requires re-evaluation of the learned reconstruction due to modifications in the leakage statistics of the same nominal grid.

Off-grid OTFS estimation has been formulated using fast Bayesian compressive sensing with a Laplace-scale-mixture prior to refine the channel representation [16]. If the path model is accurate and the components are relatively distinct, then this method can be used economically. Practical estimation becomes more difficult when multiple paths overlap within the available resolution or pulse effects depart from the assumed dictionary. Neighborhood reconstruction seeks to obtain an effective sampled response by considering the contribution of overlapping paths. Comparisons between these targets require an explicit definition of a true support detection. The physical path support score, the thresholded grid score and the cluster overlap score of the same reconstruction may be different. The grid response error evaluation method and a declared neighborhood threshold for support evaluation have been proposed. Physical path resolution requires additional verification of continuous parameter identifiability. Continuous parameter refinement can also be employed to extend the system for real-time velocity estimation in downstream applications.

2.2. Sparse Recovery and Structured Regularization

A compressed-sensing basis-expansion approach employs combined orthogonal matching pursuit to exploit segmented joint sparsity in continuous Doppler-spread OTFS channels [17]. However, a finite dictionary can approximate one fractional Doppler path with several atoms. Atom selection is used to identify true multipath and the representation error of an off-grid component. Sparse Bayesian learning has been applied to OTFS channel estimation through a probabilistic sparse-recovery formulation [18]. The performance will be affected by the chosen hierarchy, the numerical stopping condition, and the consistency of the prior with correlated leakage. A reproducible baseline specification defines the hierarchy, numerical stopping criterion and prior configuration. Computational cost also depends on dictionary size, matrix factorization reuse, and convergence tolerance, making these settings essential to interpreting iteration counts. The validation protocol therefore includes baseline tuning with a fixed observation operator and a shared validation allocation.

Meng et al. [19] exploited shared delay-Doppler-angle support in massive MIMO-OTFS and combined joint sparsity pattern learning with coefficient estimation. For fractional Doppler reconstruction, neighboring cells can likewise represent related parts of a path response; this motivates a local neighborhood model. Neighborhood geometry is a modeling choice. A fixed wide group may have noise, merge adjacent paths, and favor diffuse responses; a narrow group may also exhibit the weaknesses of point sparsity. Overlapping groups have reduced the sensitivity to arbitrary partition boundaries but introduced repeated membership and require normalization. Anisotropic neighborhoods provide a more particular prior for the discrete delay and fractional Doppler. The current structure uses integer-aligned delays as the scope condition, and fractional delays need to be added. The scope of the proposed mechanism is restricted, and thus only a small number of delay locations can be used for coherent Doppler neighborhoods. The resulting support permits several coherent coefficients per physical path. Validation of the neighborhood scale needs to be carried out for different offsets and channel densities.

2.3. Learned Reconstruction and Remaining Gaps

Time-frequency-domain learning has been used to estimate fractional delay-Doppler OTFS channels while reducing the cost of constructing a constituent channel-parameter matrix [20]. Mean squared error training can reduce noise by averaging ambiguous responses and suppressing weak components. A modified U-Net has also been investigated for channel estimation in unmanned aerial vehicle communications, providing a related encoder-decoder architecture reference [21]. The skip path can also directly transmit observation noise to the output and thus provide a reason for regulating the reconstructed response after feature fusion. Separate real and imaginary channels are used to store the complex coefficient information in a real-valued convolutional implementation, and phase equivariance is still needed. Phase augmentation and consistency checks are specified as training and validation procedures; exact global-phase equivariance of the complete real-valued network remains an architectural property to be established.

Wang et al. [22] investigated model-based deep unfolding for efficient channel estimation in zero-padded OTFS systems. The effectiveness of a model-based estimator depends on how accurately its forward operator represents the acquisition chain. Synchronization errors, imperfect pulse responses and interference outside the observation area may be present in the residual components. Environment- and signal-to-noise-ratio-adaptive channel estimation has been investigated using a two-tier adapter-switching framework [23]. The specific contribution is the coupling of residual reconstruction, a fixed anisotropic Doppler neighborhood, a noise-conditioned output threshold, and observation consistency within one effective-channel estimator. Its novelty lies in this scoped combination and its diagnostic interpretation relative to existing learned and structured sparse recovery frameworks. The evaluation distinguishes

coefficient error, support behavior, residual compatibility, and resource use. Separation is required because improvements in one area may be accompanied by declines in others, and weak path energy contributes little to the overall error of a detection but is still necessary for that detection.

3. Observation Model and Sparsity Constrained U Net

3.1. Effective Channel and Pilot Observation

Consider a single antenna frame with 64 delay cells and 32 Doppler cells. Delay indices are assumed integer aligned, while the Doppler coordinate of each path contains an integer part and a fractional offset. The channel remains constant in its delay Doppler parameters over one frame. The effective response is a sampled reconstruction target under the stated processing assumptions; physical delay resolution and practical pulse-level behavior require separate waveform validation. Each active delay location contributes a periodic Doppler leakage vector. The effective channel equation combines these contributions through their complex gains, delay locations, and Doppler coordinates:

$$H[l,k] = \sum_{p=1}^{P} a_p \delta[l - l_p] d_N(k - k_p - \kappa_p). \tag{1}$$

Here, P denotes the number of propagation paths, a_p is the complex gain of path p, and l_p identifies its delay index. The integer Doppler index k_p specifies a circular shift, while the fractional offset κ_p uses a half-open interval from minus one half, included, to one half, excluded; the integer Doppler index is adjusted when a boundary is crossed. The normalized leakage kernel d_N is obtained by averaging complex exponentials over the $N = 32$ Doppler samples. The kernel is periodic over the Doppler grid and defines a reduced effective response model with idealized pulse handling. Waveform level implementation requires phase factors and boundary effects derived from the specified transform conventions. The coefficient array has 64 rows and 32 columns, with circular boundary treatment along Doppler. Zero padding along delay preserves separation between the propagation delay boundaries. The reference gain normalization gives unit expected total channel energy before noise is added. Reproducibility requires storage of the target before visibility masking and thresholding.

The pilot observation equation relates the effective channel to the acquired samples through a known linear operator:

$$\boldsymbol{y} = \boldsymbol{A}vec(\boldsymbol{H}) + \boldsymbol{n}, \boldsymbol{n} \sim \mathcal{CN}(\boldsymbol{0}, \sigma^2 \boldsymbol{I}). \tag{2}$$

The operator A contains the transmitted pilot pattern, the acquisition mask, and the effective processing used by the receiver. Noise is modeled as circular complex Gaussian at this observation stage. Payload interference requires an explicit operator contribution or a separately characterized model mismatch term. A regularized pilot inversion produces the initial complex grid estimate $\mathbf{H}_0$. The initialization uses the received pilot observations and their known locations. The input tensor $\mathbf{X}$ concatenates the real and imaginary components of this estimate with a binary visibility mask and a broadcast noise scale. Each component is normalized using a scale obtained from the observed pilot energy and known pilot amplitude. Observation based normalization preserves the information available at inference. A single stored normalization convention is applied consistently to initialization, targets, operator action, and reconstruction metrics so that the loss terms refer to the same coefficient scale.

Figure 1 connects the pilot allocator, acquisition operator, regularized initialization, four channel input tensor, constrained reconstructor, and detector interface. The visibility mask follows a separate side branch into the network, while the known acquisition operator follows another branch into the consistency loss during training. The diagram assigns the reference channel exclusively to training supervision and derives receiver inputs from pilot observations. A stored configuration record accompanies each observation and identifies grid dimensions, pilot energy, pulse convention, and noise normalization. The detector interface receives the effective reconstructed grid together with the boundary convention used by the forward operator. These interfaces define the estimator together with its network weights; changes to the acquisition or detector convention therefore require reassessment of reconstruction behavior.

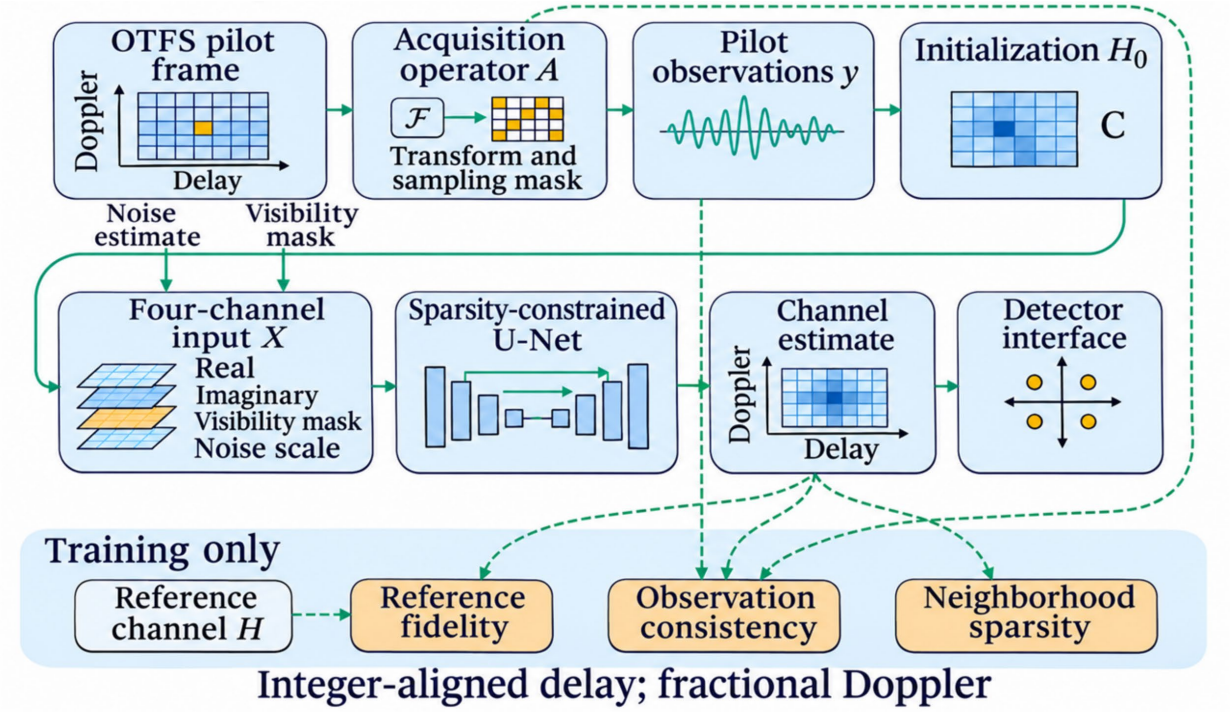


Figure 1. Pilot observation and reconstruction workflow from the transmitted frame to the effective delay-Doppler channel estimate

3.2. Residual Encoder Decoder and Neighborhood Shrinkage

The encoder has three resolutions with feature widths of 24, 48, and 96 channels. At full resolution, two padded convolutions use kernels of three by five cells to emphasize Doppler context while retaining delay information. A stride of two along both axes produces intermediate grids of 32 by 16 and 16 by 8. The bottleneck contains two residual convolutions at 96 channels. Bilinear interpolation followed by convolution restores resolution in the decoder, and concatenated skip connections supply the corresponding encoder features. Group normalization reduces dependence on batch statistics, while smooth nonlinear activation supports stable gradients near weak responses. A final linear convolution predicts two residual channels. Their complex representation is added to the initial estimate before the constraint stage. This residual design preserves reliable observation content while directing network capacity toward noise, missing cells, and systematic leakage errors.

To regulate the reconstructed support, the neighborhood shrinkage equation scales each complex coefficient according to the energy in its surrounding Doppler cells:

$$\boldsymbol{Z} = \boldsymbol{H}_0 + F_\theta(\boldsymbol{X}),$$
$$\widehat{\boldsymbol{H}} = \boldsymbol{Z} \odot max\left(0{,}1 - \frac{\tau}{\sqrt{G(|\boldsymbol{Z}|^2)} + \varepsilon}\right). \quad (3)$$

The neighborhood operator G averages squared magnitude over a one by five Doppler window with circular padding and unit sum weights. The positive threshold τ is conditioned on the noise scale through a small learned affine map followed by a positivity enforcing activation. The stabilizer ε prevents division by zero and is fixed independently of channel magnitude. The maximum and scaling operations are applied elementwise, and $\odot$ denotes elementwise multiplication. The shrinkage map is piecewise differentiable, with a change in derivative at the activation boundary. Shrinkage multiplies the real and imaginary components by the same real factor, preserving their local complex phase wherever the factor is nonzero. Local energy sharing permits retention of a weak coefficient within an energetic neighborhood. Phase preservation applies to the shrinkage stage; attenuation can still bias coefficient amplitudes and does not establish phase equivariance of the full network. It can also retain noise close to a strong path; the consistency loss and held out threshold calibration address this effect. Fixed spatial weights make the support width inspectable and preserve the explicit neighborhood interpretation of the constraint.

Figure 2 places the three encoder scales on the left, the bottleneck centrally, and the mirrored decoder scales on the right. Resolution annotations identify feature geometry, while separate arrows indicate the initial estimate bypass and the noise conditioning path. The output stage first combines the learned residual with the pilot initialization and then applies neighborhood shrinkage. During training, the resulting estimate is connected to both a reference reconstruction loss and the known observation operation. During inference, the loss branches are removed. At deployment, all inputs

are derived from the received pilot and known acquisition settings. Parameter count and computational latency depend on the selected tensor framework, padding convention, and hardware precision. Hardware timing therefore requires a concrete implementation and a specified profiling procedure.

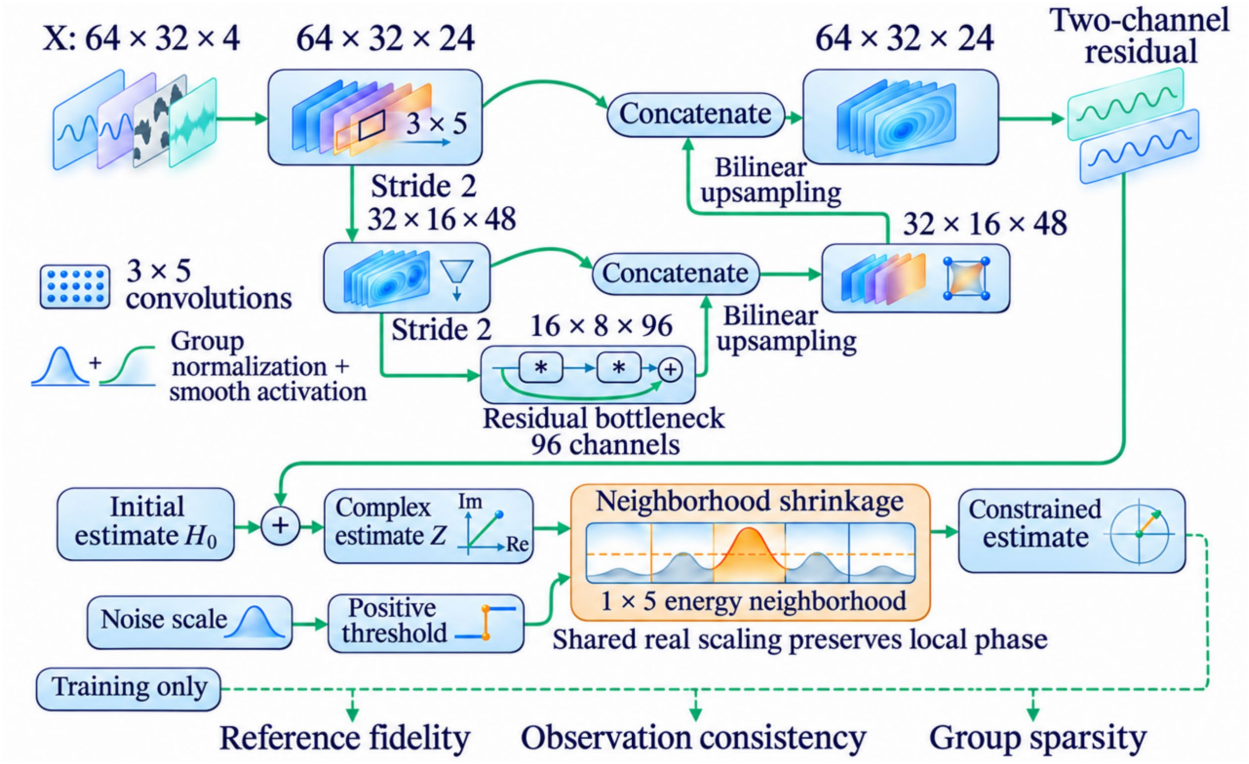


Figure 2. Residual U-Net architecture with anisotropic convolutions, skip connections, noise-conditioned neighborhood shrinkage, and training losses

3.3. Learning Objective and Implementation Protocol

The training objective balances reference fidelity, compatibility with the received pilot, and neighborhood sparsity. Reference fidelity is normalized per frame to prevent energetic channels from dominating optimization. Observation consistency is normalized by the number of observed pilot samples and their noise variance. Group penalization acts directly on the reconstructed response to retain an interpretable output support. These contributions enter the learning objective equation with separate weights for observation consistency and group sparsity:

$$\mathcal{L} = \mathbb{E}\left[\frac{\|\widehat{\boldsymbol{H}} - \boldsymbol{H}\|_F^2}{\|\boldsymbol{H}\|_F^2 + \varepsilon} + \lambda_c \frac{\|\boldsymbol{A}vec(\widehat{\boldsymbol{H}}) - \boldsymbol{y}\|_2^2}{m\sigma^2} + \lambda_g \sum_{g \subset \mathcal{G}} \|\widehat{\boldsymbol{H}}_g\|_2\right]. \tag{4}$$

The expectation is taken over training frames, with independent validation used for loss selection and a held-out partition reserved for testing. The set $\mathcal{G}$ contains overlapping output neighborhoods, and $\widehat{\mathbf{H}}_g$ denotes the membership normalized coefficient vector associated with group g. These groups use the same one by five geometries as the shrinkage operator, with normalization controlling repeated coverage. Each coefficient is divided by the square root of its neighborhood membership count before a group vector is formed; with all centered one-by-five circular Doppler windows, that count is five. The number of observed pilot samples is denoted by m. Initial validation candidates are 0.1 for the consistency coefficient λ_c and 0.002 for the group coefficient λ_g. Both require retuning when operator normalization changes.

The following settings specify the proposed training protocol, whose execution requires retained checkpoints and validation traces. The implementation specification uses Adam optimization with an initial learning rate of 0.001, batches of 32, and a maximum of 150 epochs. Early stopping monitors validation coefficient error with a patience of 15 epochs. Random global phase rotation is applied jointly to inputs and targets. Noise and Doppler offsets are resampled within each training split, while physical path seeds remain separated across splits. Gradient clipping and a positive lower bound on the noise estimate reduce instability associated with low energy frames.

A reproducible implementation retains the exact forward operator, split identifiers, optimizer state, selected checkpoint, and validation trace. Three matched ablations remove neighborhood penalization, observation consistency, or both while retaining the residual architecture and training allocation. A point shrinkage variant replaces neighborhood energy with coefficient magnitude while preserving the remaining network. These tests distinguish the architectural effect from the effect of increased parameter capacity. The constraint is evaluated against three failure criteria: loss of the error advantage under matched tuning, disproportionate deterioration in weak path recall, and increased residual incompatibility outside the training range. The numerical analysis follows the metric definitions and component comparisons specified by this procedure.

4. Numerical Evaluation and Error Analysis

4.1. Evaluation Protocol and Nominal Reconstruction

Nominal reconstruction behavior, component sensitivity, and robustness are examined under a common evaluation protocol. Analysis of the coefficient error is related to support recovery and observation residuals. A unified observation geometry and normalization convention are used to conduct the comparison of estimator configurations. Gaussian-sinc pulse shaping for Zak-OTFS illustrates how pulse design enters waveform-specific channel modeling and input-output relation estimation [24]. The evaluation configuration uses a 64 by 32 effective grid, integer delays in the first 16 cells, six nominal paths, and fractional Doppler offsets covering the half-bin interval. Superimposed full-pilot schemes have been studied with iterative channel estimation for multiuser OTFS uplinks, representing a distinct pilot-allocation strategy [25]. The present comparison uses an embedded-pilot observation with a nominal observation fraction of 12.5%, and a shared pilot energy convention is used for the five estimator labels: threshold initialization, orthogonal matching pursuit, sparse Bayesian learning, unconstrained U-Net, and the constrained model. Table 1 lists the assumptions and the proposed training, validation and test divisions. Physical path seeds are divided before noise generation, and observations obtained from a single physical realization belong to that split. Thus, the separation of channel familiarity and noise diversity can be achieved. The observation fraction denotes the sampled share of the effective grid. Pilot energy and acquisition geometry are held fixed across configurations at each comparison point, while observation-fraction sweeps change only the declared sampling allocation.

Table 1. Common evaluation protocol and parameter roles

Item	Setting	Parameter role
Grid and path alignment	64 delay × 32 Doppler; integer delay	Channel geometry
Nominal paths and offset	6 paths; offset within half bin	Path configuration
Observation fraction	12.5 percent nominal	Observation setting
Planned split	40000 / 5000 / 10000 frames	Partition allocation
SNR sweep	0 to 30 dB in 5 dB increments	Evaluation coordinates
Support rule	−20 dB relative to peak	Declared operating point
Runtime values	Batch one processing	Computational allocation

The nominal signal to noise ratio is defined through the average noise free pilot observation energy and the average complex noise energy. Reconstruction error covers the entire effective response, including cells outside the visible acquisition mask. Framewise normalized errors are averaged in linear units before conversion to decibels so that the aggregate metric retains its energy interpretation. The support protocol thresholds each response at minus 20 dB relative to its own frame peak and uses one-to-one matching at the same delay with a one-cell circular Doppler tolerance. Precision counts matched detections among accepted detections; recall counts matched reference cells among reference support cells. Empty-support cases require a separately reported convention. Deep-learning-assisted path prediction and channel estimation for high-delay-resolution MIMO-OTFS provide a related example of explicitly modeling propagation-path information [26]. The support threshold and matching tolerance used here define the present evaluation protocol. Figure 3(a) displays the error sweep at signal to noise ratios from 0 to 30 dB in 5 dB increments. The constrained values are −3.5, −7.0, −12.0, −17.5, −22.4, −25.0, and −26.0 dB. Relative to the unconstrained values of −18.8 dB at 20 dB and −21.0 dB at 30 dB, the corresponding separations are 3.6 and 5.0 dB. The constrained curve improves by 4.9 dB between 15 and 20 dB but by only 1.0 dB between 25 and 30 dB. This flattening is consistent with an increasing contribution of representation bias as observation noise decreases.

Variational Bayesian estimation for MIMO OTFS/OTSM with finite-resolution converters explicitly incorporates quantization into the observation model [27]. The present distribution analysis uses a fixed acquisition model and separates typical coefficient error from the spread of framewise errors. The 160 displayed values per configuration

summarize the plotted collection, whereas the 10,000-frame test allocation in Table 1 belongs to the proposed protocol; independent training repetitions require separate records. Figure 3(b) presents cumulative probability points generated from 160 samples for each label, with distribution centers of −10.8, −14.2, −16.8, −18.8, and −22.4 dB and a common 1.4 dB sampling scale. The constrained and unconstrained locations differ by 3.6 dB, whereas their distribution widths are comparable. This arrangement separates a shift in typical error from a change in variability. The upper probability region is especially relevant to the future evaluation because a favorable average can coexist with occasional large errors. Paired channels will permit those failures to be attributed to path density, relative gain, or neighborhood overlap. Figure 3(c) examines offset magnitudes of 0, 0.1, 0.2, 0.3, 0.4, and 0.5. Constrained error changes from −23.4 to −20.8 dB, a 2.6 dB deterioration; unconstrained error changes from −20.6 to −16.9 dB, a 3.7 dB deterioration. The separation consequently widens from 2.8 to 3.9 dB. Along this trajectory, the neighborhood constraint expresses greater tolerance to leakage broadening, with its residual degradation identifying the need to examine how the five-cell window covers distant Doppler tails.

Figure 3(d) provides a leakage error surface normalized to a maximum of 0 dB and displayed over a range extending to −40 dB. The main concentration lies at delay bin 6 and zero relative Doppler, with a weaker component near delay bin 11 and relative Doppler bin 3. Their narrow delay extent and broader Doppler extent reflect the anisotropic geometry assumed by the proposed constraint. A one by five neighborhood centered on the dominant component covers relative Doppler bins −2 through 2, while more distant tails fall outside that local energy calculation. The central region supplies a support anchor, while the outer region provides a diagnostic of truncation bias. Independent cell selection represents the displayed structure through several active entries; neighborhood grouping shares evidence along the leakage direction. The two delay locations remain separate, allowing the constraint to preserve delay discrimination. Residual concentration around a correctly localized path motivates inspection of shrinkage strength and neighborhood width. A broadly distributed residual motivates examination of observation noise and operator mismatch. Peak normalization makes the map a spatial diagnostic at one displayed configuration. Its color range describes relative error within that map; comparisons of absolute accuracy across configurations require the energy-based metric and a common normalization.

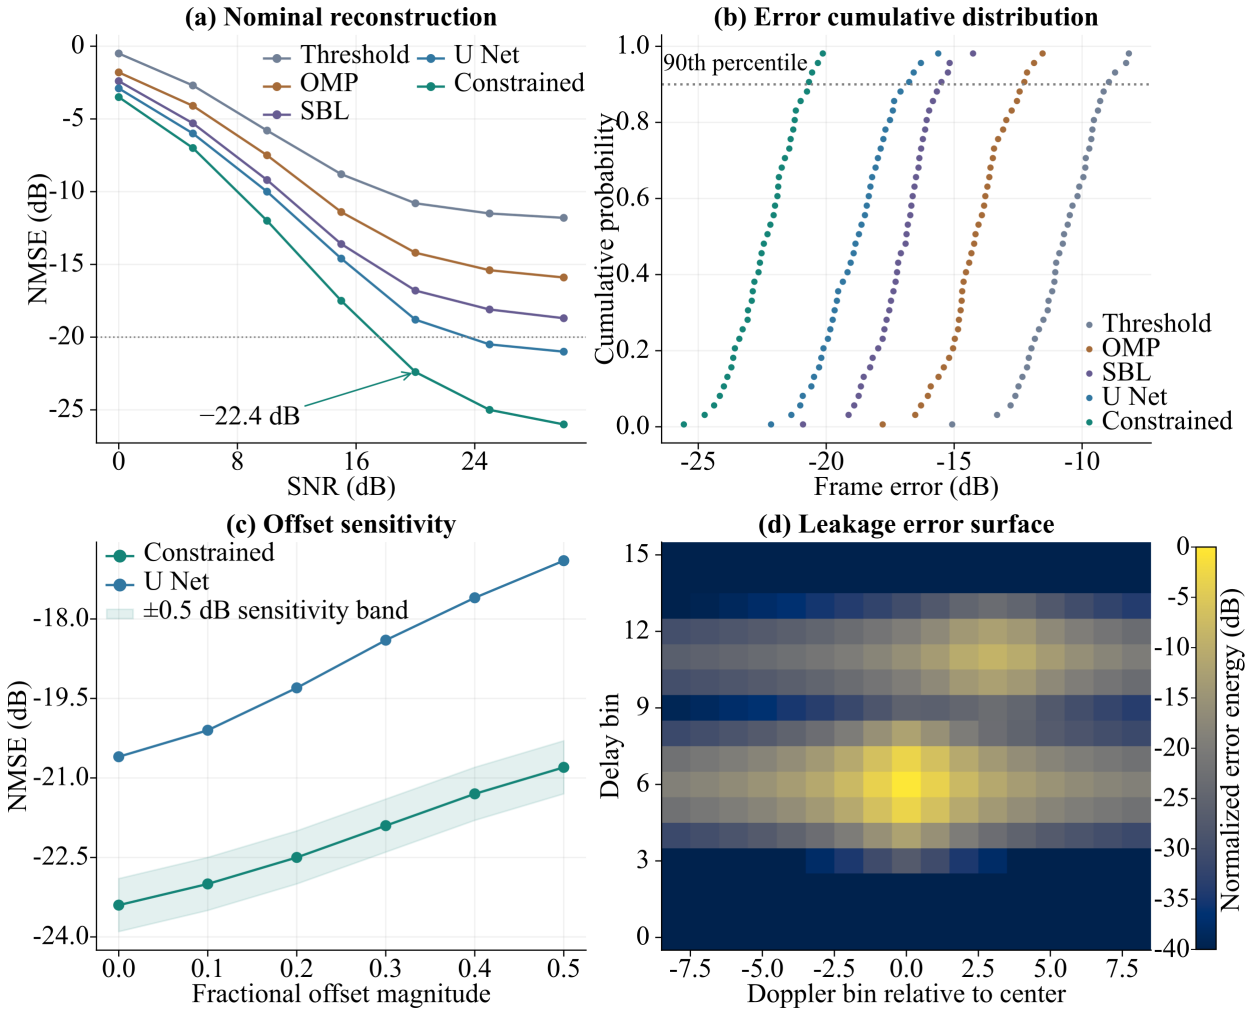


Figure 3. Nominal reconstruction analysis: (a) error versus signal-to-noise ratio; (b) error cumulative distributions; (c) fractional-offset sensitivity; (d) normalized leakage error map

4.2. Constraint Ablation and Support Sensitivity

Sparse-prior-guided deep learning has been investigated for OTFS channel estimation, placing an explicit sparsity prior within a learned reconstruction framework [28]. The present ablation holds optimization and validation resources fixed across variants. The planned variants share encoder widths, input channels, data partitions, initialization rules, and early stopping criteria. The comparison record must also identify trainable parameter counts and equal validation-search budgets before attributing differences to a constraint. Their remaining loss weights are calibrated against the same validation allocation. This protocol separates the intended constraint comparison from the effects of model capacity and unequal tuning. Figure 4(a) locates five ablation values: −18.8 dB for the unconstrained variant, −20.1 dB for consistency alone, −20.9 dB for sparsity alone, −21.3 dB for point shrinkage, and −22.4 dB for the complete neighborhood model. Each interval extends 0.6 dB on either side of its target and denotes a prescribed sensitivity band, distinct from a confidence interval estimated from repeated training runs. Relative to the unconstrained target, consistency and sparsity contribute error differences of 1.3 and 2.1 dB, respectively. The complete model is separated by 3.6 dB, while neighborhood sharing contributes a further 1.1 dB relative to point shrinkage. These comparisons indicate complementary roles for observation compatibility and neighborhood support. Paired frame comparisons can further resolve how the joint constraint changes the channel realizations that benefit. The intervals help identify comparison scales that the future experiment should resolve.

Figure 4(b) displays 160 error samples for each ablation variant, centered on the five values with a common 0.9 dB sampling scale. The full model distribution is displaced toward smaller errors, while overlapping distributions preserve the possibility of unfavorable individual realizations. Such overlap matters for interpreting a regularizer: improving the median and preserving a narrow upper tail are distinct objectives. A future paired analysis should identify whether weak components, nearly colliding paths, or unusual gain ratios populate that upper tail. Grid-evolution methods for doubly fractional OTFS estimation adapt an initially uniform grid to improve the trade-off between representation resolution and recovery complexity [29]. In the present fixed-grid model, coefficient attenuation and neighborhood selection are examined separately. The ablation therefore requires separate inspection of errors inside accepted neighborhoods and errors outside them. Strong attenuation within a correctly detected neighborhood would point toward amplitude bias; substantial energy outside accepted neighborhoods would indicate missed leakage or unsupported activity. The point shrinkage comparison isolates local energy sharing by retaining the residual architecture and applying independent coefficient suppression. An advantage concentrated at larger offsets is consistent with the proposed neighborhood geometry. Deterioration near integer alignment motivates noise conditioned adjustment of the neighborhood width. The distributions expose the error range over which amplitude bias and support omissions require separate examination.

Figure 4(c) separates support precision from recall for the five estimator labels. The coordinate pairs are 0.62 and 0.71 for threshold initialization, 0.74 and 0.77 for pursuit, 0.79 and 0.83 for Bayesian recovery, 0.82 and 0.86 for unconstrained U-Net, and 0.91 and 0.89 for the constrained label, with precision stated first. The last comparison increases precision by 0.09 and recall by 0.03. Expressed as complementary fractions, the unsupported share among accepted detections decreases from 0.18 to 0.09, while the missed share of reference support decreases from 0.14 to 0.11. The two fractions quantify distinct conditional proportions because their denominators differ. Most of the support improvement corresponds to suppression of unsupported activity, accompanied by a smaller improvement in coverage. This behavior is relevant to a neighborhood constraint because local energy sharing can preserve a weak coefficient adjacent to a coherent response while suppressing isolated activity elsewhere. The same mechanism can also preserve noise close to a dominant path. Stratification by distance from a strong component distinguishes local contamination from recovery of weak support. Gain stratification is also necessary: an apparently favorable overall recall can still conceal poor recovery of the weakest path population.

Figure 4(d) examines thresholds of −28, −24, −20, −16, and −12 dB. Precision increases through 0.77, 0.84, 0.91, 0.95, and 0.98, while recall decreases through 0.97, 0.94, 0.89, 0.80, and 0.65. Raising the threshold from −24 to −20 dB increases precision by 0.07 while reducing recall by 0.05. A further increase from −20 to −16 dB gains 0.04 precision but sacrifices 0.09 recall, and the final increment gains only 0.03 precision while sacrificing 0.15 recall. The marginal tradeoff therefore becomes progressively less favorable for applications that depend on weak path coverage. The nominal −20 dB setting retains the 0.91 precision and 0.89 recall balance. A threshold relative to frame peak remains sensitive to dominant path gain, so its calibration should be compared with a noise normalized alternative. Freezing the selected rule before test evaluation preserves the separation between calibration and assessment. The displayed progression directs attention toward frames with large dynamic range: a threshold that suppresses an isolated noise cell in one frame may suppress a meaningful weak tail in another. Framewise assessment can resolve the relative influence of acquisition reliability and dominant coefficient magnitude on adaptive shrinkage.

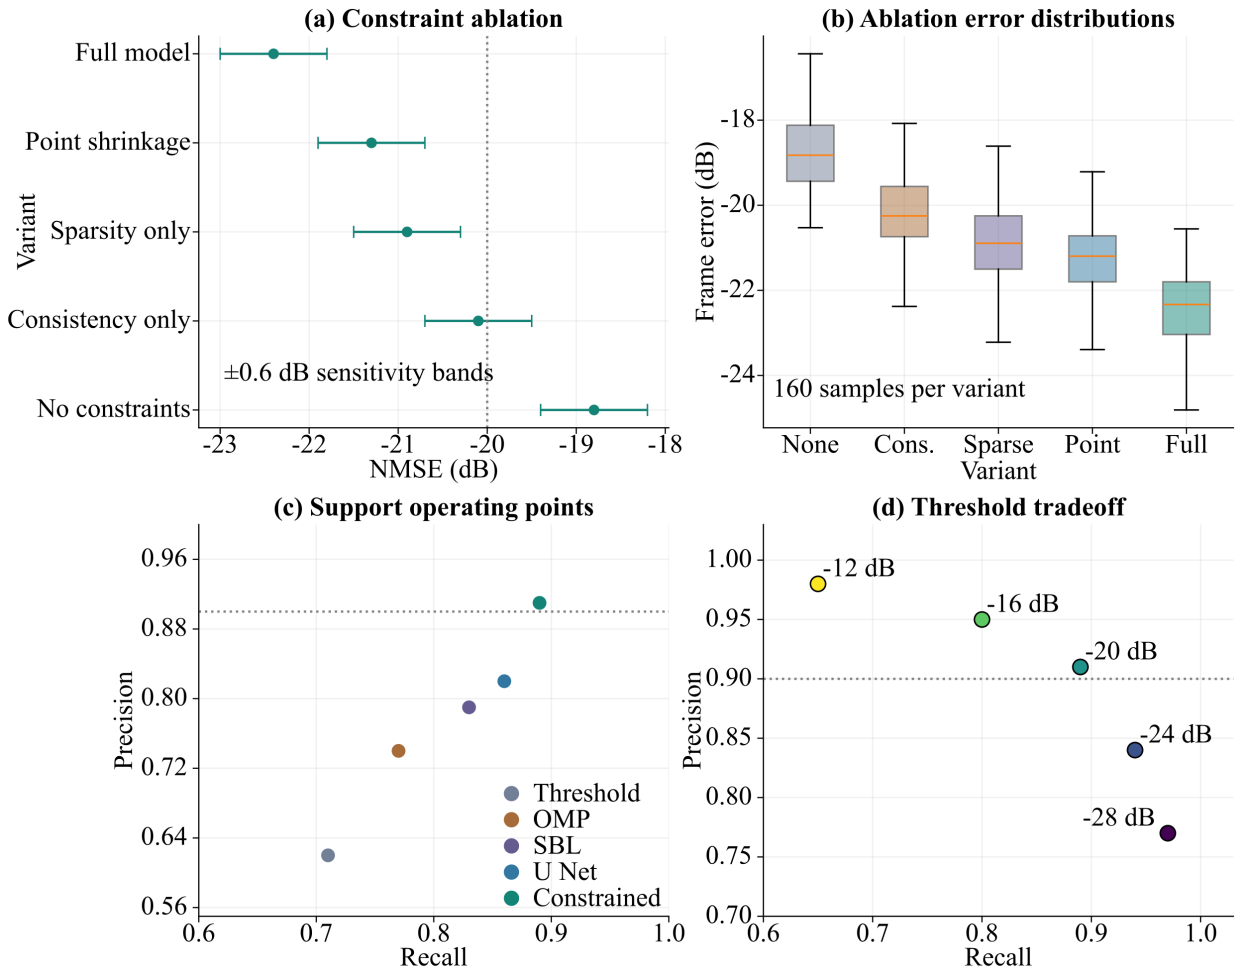


Figure 4. Constraint and support analysis: (a) ablation error sensitivity bands; (b) ablation error distributions; (c) precision-recall operating points; (d) threshold-dependent support tradeoff

4.3. Robustness and Computational Diagnostics

Representation-based continual learning has been investigated for channel estimation in dynamic wireless environments, addressing adaptation as propagation conditions change [30]. The robustness analysis therefore varies observation fraction, path count, offset magnitude, and noise scale separately. Figure 5(a) shows constrained errors of −17.0, −20.1, −22.4, −23.6, and −24.0 dB at observation fractions of 6.25, 9.375, 12.5, 18.75, and 25 percent. The unconstrained counterparts are −14.8, −17.2, −18.8, −20.0, and −20.6 dB. Increasing the fraction from 6.25 to 12.5 percent improves the constrained target by 5.4 dB, whereas doubling it again improves the target by 1.6 dB, expressing diminishing returns from additional observations. Figure 5(b) shows a deterioration from −25.1 to −17.6 dB as path count rises from 2 to 12; the unconstrained targets change from −22.0 to −13.4 dB. Greater density increases neighborhood overlap and weakens the distinction between compact leakage and diffuse activity. Figure 5(c) gives penalties of 2.8, 1.1, 0, 0.2, 0.9, and 2.4 dB for noise scale mismatches of −6, −3, 0, 1, 3, and 6 dB. The mismatch sign describes the supplied noise scale relative to its matched value: negative values indicate underestimation and positive values indicate overestimation. The reported penalties are error differences relative to the matched setting; the amplitude or power convention must accompany an executable configuration. Figure 5(d) shows offset conditioned distributions centered on −23.4, −23.0, −22.5, −21.9, −21.3, and −20.8 dB. Their increasing sampling scales expose the joint effect of a deteriorating location and widening variability.

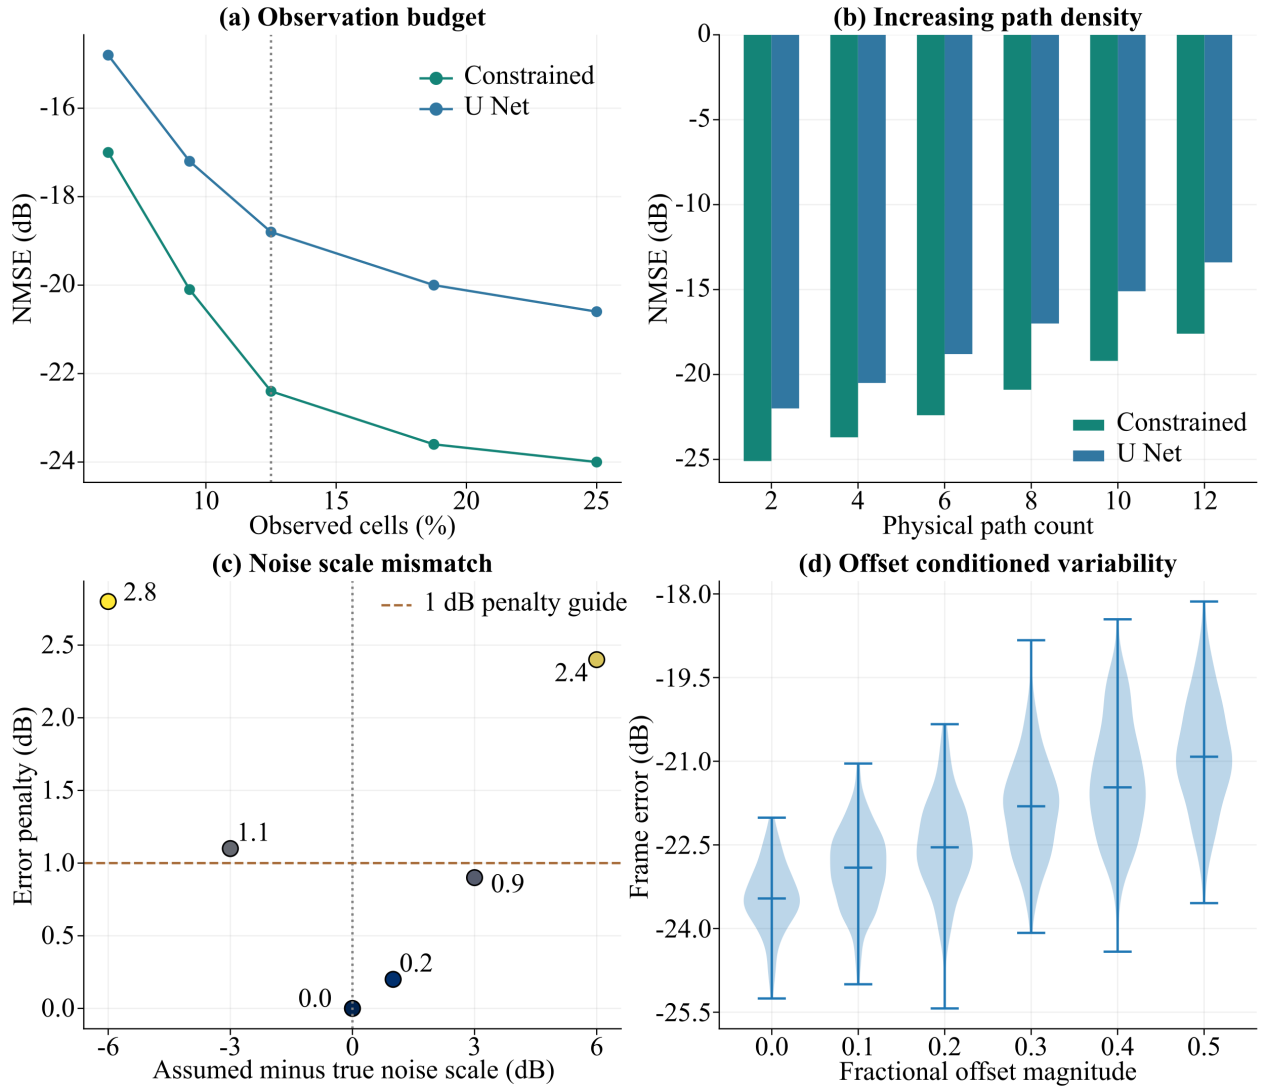


Figure 5. Robustness analysis: (a) observation fraction response; (b) path density comparison; (c) noise scale mismatch penalties; (d) offset-conditioned violin distributions

Fast burst-sparsity learning for massive MIMO-OTFS combines a hybrid sparse prior with variational inference to reduce the cost of high-dimensional recovery [31]. For the present computational comparison, batch size, numerical precision, synchronization, and memory-transfer policy are reporting conditions. The millisecond coordinates are framewise computational budgets. Device-specific latency requires an executed implementation, identified hardware and software, warm-up, synchronized repeated timing, and variability reporting. Figure 6(a) pairs nominal errors with budgets of 0.08 ms for threshold initialization, 2.6 ms for pursuit, 7.8 ms for Bayesian recovery, 0.65 ms for unconstrained U-Net, and 0.74 ms for the constrained model. The −20 dB and 1 ms guides define a joint design target. The constrained point lies within the region defined jointly by the two guides; the remaining configurations exceed at least one guide. Its budget exceeds the unconstrained allocation by 0.09 ms, or approximately 13.8 percent, while its error improves by 3.6 dB. This ratio quantifies the accuracy and computational tradeoff at the specified operating point. Figure 6(b) allocates 0.08, 0.17, 0.12, 0.28, and 0.09 ms to initialization, encoding, bottleneck processing, decoding, and shrinkage. Decoding accounts for approximately 37.8 percent of the total and shrinkage for 12.2 percent. The decoder allocation motivates profiling interpolation and concatenated skip features, whose memory traffic may become consequential at larger grid dimensions. Comparable timing requires identical initialization and preprocessing for both neural variants. Batch one latency, sustained throughput, parameter storage, and peak activation memory characterize separate implementation requirements. The component breakdown identifies the operations whose cost should be examined as the grid dimensions increase.

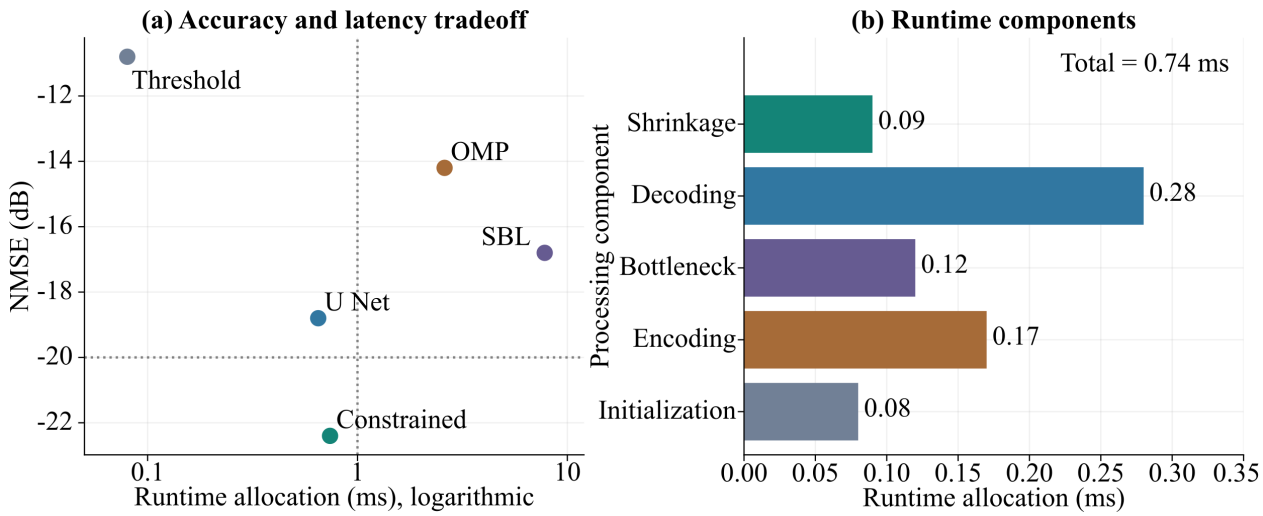


Figure 6. Computational allocations: (a) computational budget and reconstruction error trade-off; (b) constrained estimator budget component breakdown

Parametric bilinear inference for OTFS jointly estimates channel gains and data symbols using Gaussian belief propagation, linking channel recovery with the demodulation model [32]. The present residual diagnostic examines compatibility between a reconstructed response and its specified pilot observation. Figure 7(a) shows residual trajectories across 12 checkpoints: the unconstrained response decreases from 2.4 to 1.36, and the constrained response decreases from 2.5 to 1.02. At the final checkpoint, the constrained excess above the unit guide is 0.02 compared with 0.36 for the unconstrained label. The two trajectories approach the guide at different rates, with the constrained sequence reaching 1.10 at checkpoint 7 and 1.05 at checkpoint 9. This progression indicates closer agreement with the observation normalization. The unit guide remains conditional on the stated noise model, regularization, and effective degrees of freedom. Proximity to this guide is a compatibility diagnostic; statistical calibration and channel-recovery accuracy require their own assessments. Correlation between estimated coefficients and observation noise can shift the expected residual, while unmodeled interference can create systematic excess. A future implementation should therefore pair residual energy with residual structure, inspecting whether remaining activity is concentrated around paths or distributed throughout the observation region. A localized residual can indicate attenuation of a physically coherent component; a broader residual can reflect acquisition mismatch or noise calibration. Tracking this distinction alongside coefficient error prevents a low residual from dominating interpretation when the observation operator leaves substantial ambiguity in the reconstructed grid.

Model-driven OTFS detectors have been designed to account for channel estimation error, providing a direct motivation for evaluating reconstruction through a specified detector [33]. Figure 7(b) shows retained core fractions of 78, 78, 77, 74, and 68 percent at relative thresholds of −28, −24, −20, −16, and −12 dB. Retained tail fractions are 18, 17, 15, 12, and 8 percent, while removed fractions are 4, 5, 8, 14, and 24 percent. Each composition sums to 100 percent. Moving from −24 to −20 dB removes another 3 percentage points of energy, including 1 point from the core and 2 points from the tail. Moving from −20 to −12 dB removes another 16 points, with core retention falling by 9 points and tail retention by 7 points. Tail suppression therefore becomes increasingly accompanied by attenuation of the central response. The nominal setting retains 92 percent of the reference energy, compared with 76 percent at the highest threshold. The retained-energy composition is a channel-domain diagnostic. A fixed-detector study should examine its relationship to bit error rate under the same observation and equalization settings. Fractional delay, synchronization offsets, and overlapping paths should be introduced through separate perturbations, with a matched clean reference retained for each case. Their energy compositions can distinguish support truncation, coefficient attenuation, and changes in the effective observation operator.

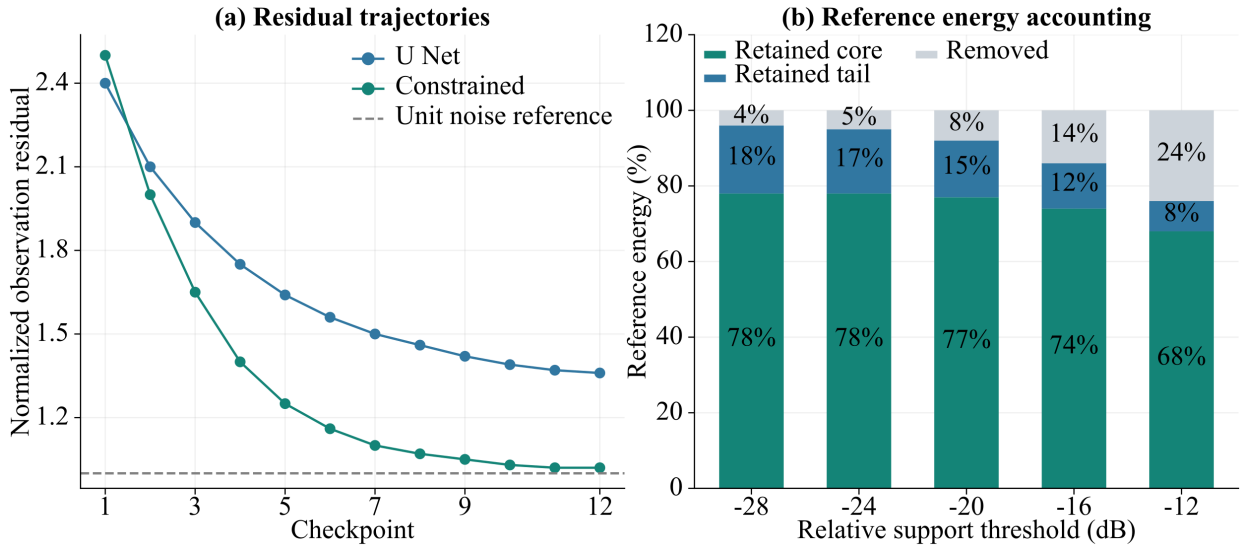


Figure 7. Diagnostic analysis: (a) normalized observation residual versus checkpoint; (b) retained core, retained tail, and removed energy across support thresholds

5. Conclusion

This study develops a delay-Doppler sparsity-constrained U-Net and uses the reported numerical coordinates to examine the association between neighborhood reconstruction, coefficient accuracy, and support recovery under the stated effective-channel assumptions.

1. Reconstruction accuracy. At a signal to noise ratio of 20 dB, the constrained configuration has a reported normalized mean square error of −22.4 dB, compared with −18.8 dB for the unconstrained U-Net, corresponding to a 3.6 dB improvement. At 30 dB, this difference increased to 5.0 dB. The complete model also improved upon the point shrinkage variant by 1.1 dB, supporting the contribution of shared neighborhood energy to fractional Doppler leakage reconstruction.

2. Support recovery and offset sensitivity. Support precision increased from 0.82 to 0.91, while recall increased from 0.86 to 0.89. As fractional Doppler offset magnitude increased from zero to 0.5, reconstruction error deteriorated by 2.6 dB for the constrained model and 3.7 dB for the unconstrained model. These patterns associate neighborhood shrinkage with stronger suppression of unsupported activity and greater tolerance to leakage broadening.

3. Contribution and applicability. Residual reconstruction, noise-conditioned neighborhood shrinkage and observation consistency are used to construct a structure in an interpretable estimation framework. Its applicability depends on integer-aligned delays, a known acquisition operator, and appropriate noise calibration. A fixed neighborhood can merge closely spaced responses or retain noise near a dominant path. Fractional delay can spread energy across delay cells, while synchronization and pulse mismatch alter the observation operator; these extensions require separate validation.

Future work should investigate adaptive neighborhood selection, fractional-delay processing, and joint evaluation with symbol detection under practical receiver conditions.

6. Declarations

6.1 Author Contributions

Riccardo Amato contributes to conceptualization, methodology, software, validation, analysis, investigation, data collection, draft preparation, manuscript editing, visualization, supervision, project administration, and funding acquisition. Stefano Di Marco contributes to validation, analysis, investigation, data collection, draft preparation. Antonio Ferrara contributes to conceptualization, validation, analysis. Emiliano Gentile contributes to analysis, investigation, data collection. All authors have read and agreed with the manuscript before its submission and publication.

6.2 Funding and Acknowledgments

This research received no specific financial support from any funding agency.

6.3 Institutional Review Board Statement

Not applicable.

6.4 Informed Consent Statement

Not applicable.

6.5 Date Availability

Not applicable.

6.6 Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

6.7 Declaration of generative AI and AI-assisted technologies

During the preparation of this manuscript, the authors utilized Grammarly for linguistic polishing, logical organization and draft revision. The generative AI tool was not involved in research design, data collection, data analysis, or the equation of core conclusions. All outputs generated by AI were reviewed, revised and verified manually by the authors. The authors take full responsibility for all content of this manuscript.

7. References

[1] Aldababsa, M., Özyurt, S., Kurt, G. K., & Kucur, O. (2024). A survey on orthogonal time frequency space modulation. IEEE Open Journal of the Communications Society, 5, 4483–4518. https://doi.org/10.1109/ojcoms.2024.3422801

[2] Sheng, H.-T., & Wu, W.-R. (2024). Time-frequency domain channel estimation for OTFS systems. IEEE Transactions on Wireless Communications, 23(2), 937–948. https://doi.org/10.1109/twc.2023.3283578

[3] Jesbin, F., & Chockalingam, A. (2026). Delay-Doppler pulse shaping in Zak-OTFS using Hermite basis functions. IEEE Transactions on Vehicular Technology, 75(10), 23804–23817. https://doi.org/10.1109/tvt.2026.3703829

[4] Liu, C., Xiang, W., Long, H., & Jia, Y. (2024). Spectral-efficient reference signal design and iterative channel estimation for OTFS modulation. IEEE Transactions on Vehicular Technology, 73(10), 15286–15301. https://doi.org/10.1109/tvt.2024.3414468

[5] Wang, X., Zheng, C., Hu, P., Yang, J., & Kang, C. G. (2025). A sparsity-agnostic SL0 channel estimation approach for OTFS systems. IEEE Communications Letters, 29(5), 1097–1101. https://doi.org/10.1109/lcomm.2025.3554379

[6] Zegrar, S. E., Sümer, A. S., & Arslan, H. (2025). Fractional delay and fractional Doppler estimation and mitigation framework in OTFS systems. IEEE Transactions on Vehicular Technology, 74(2), 2884–2896. https://doi.org/10.1109/tvt.2024.3470839

[7] Lee, H.-G., Kim, J., Zhao, J., & Joung, J. (2025). Single-tone-based channel estimation method for OTFS systems. IEEE Transactions on Communications, 73(3), 1638–1651. https://doi.org/10.1109/tcomm.2024.3455240

[8] Li, W., Zhou, D., & Ye, Q. (2025). Channel computation based on multi-scale attention residual network. Journal of Information and Intelligence, 3(3), 275–287. https://doi.org/10.1016/j.jiixd.2025.03.001

[9] Al-Imran, Nazim, M. S., Nguyen, H., & Jang, Y. M. (2025). Channel estimation of massive MIMO FSO communication system using deep attention residual U-Net. ICT Express, 11(2), 287–292. https://doi.org/10.1016/j.icte.2024.09.012

[10] Wang, Q., Chen, X., Tao, Q., Qian, L. P., Kam, P.-Y., & Wu, Y. (2025). Model-driven channel estimation network for orthogonal time-frequency space systems. IEEE Transactions on Vehicular Technology, 74(7), 11552–11556. https://doi.org/10.1109/tvt.2025.3548639

[11] He, Y., & Joseph, G. (2025). Bayesian algorithms for Kronecker-structured sparse vector recovery with application to IRS-MIMO channel estimation. IEEE Transactions on Signal Processing, 73, 142–157. https://doi.org/10.1109/tsp.2024.3416543

[12] Luan, D., & Thompson, J. S. (2026). Achieving robust channel estimation neural networks by designed training data. IEEE Transactions on Cognitive Communications and Networking, 12, 2311–2324. https://doi.org/10.1109/tccn.2025.3592299

[13] Zhang, H., Huang, X., & Zhang, J. A. (2024). Low-overhead OTFS transmission with frequency or time domain channel estimation. IEEE Transactions on Vehicular Technology, 73(1), 799–811. https://doi.org/10.1109/tvt.2023.3305921

[14] Wang, R., Ren, H., Pan, C., Weng, R., & Wang, J. (2026). A cyclic shift embedded pilot-based channel estimation for multiuser MIMO-OTFS systems with fractional delay and Doppler. IEEE Transactions on Signal Processing, 74, 4466–4482. https://doi.org/10.1109/tsp.2026.3721735

[15] He, X., Yuan, W., & Fan, P. (2024). On the pilot-aided channel estimation for windowed OTFS with data interference in rapidly time-varying channels. IEEE Transactions on Wireless Communications, 23(11), 16359–16374. https://doi.org/10.1109/twc.2024.3440586

[16] Tang, M., Wang, H., Yuan, Z., & Yuan, J. (2024). A novel off-grid channel estimation with fast BCS using LSM prior for OTFS modulation. IEEE Transactions on Wireless Communications, 23(9), 12157–12171. https://doi.org/10.1109/twc.2024.3388449

[17] Guo, Q., Jiang, H., Xiang, J., & Zhong, Y. (2024). A CS-BEM OTFS channel estimation approach for sparse continuous Doppler-spread channels. IEEE Wireless Communications Letters, 13(11), 2985–2989. https://doi.org/10.1109/lwc.2024.3432090

[18] Lan, X., Zhang, X., Zhang, L., Ma, Z., Wang, Q., & Akan, O. B. (2026). A sparse Bayesian learning-based OTFS channel estimation for 6G wireless system. IEEE Wireless Communications Letters, 15, 805–809. https://doi.org/10.1109/lwc.2025.3642986

[19] Meng, K., Yang, S., Wang, X.-Y., Bu, Y., Tang, Y., Zhang, J., & Hanzo, L. (2024). Joint sparsity pattern learning based channel estimation for massive MIMO-OTFS systems. IEEE Transactions on Vehicular Technology, 73(8), 12189–12194. https://doi.org/10.1109/tvt.2024.3375027

[20] Mattu, S. R., & Chockalingam, A. (2024). Learning in time-frequency domain for fractional delay-Doppler channel estimation in OTFS. IEEE Wireless Communications Letters, 13(5), 1245–1249. https://doi.org/10.1109/lwc.2024.3367112

[21] Gupta, C., & Yadav, S. S. (2025). Deep learning based channel estimation for UAVs: A modified U-Net approach. Advances in Electrical and Computer Engineering, 25(1), 61–70. https://doi.org/10.4316/aece.2025.01007

[22] Wang, X., Zhang, H., Yang, Q., Wang, B., Gui, G., Yuen, C., & Sari, H. (2026). Model-based deep unfolding for efficient zero-padded OTFS channel estimation. IEEE Transactions on Vehicular Technology, 75(1), 756–768. https://doi.org/10.1109/tvt.2025.3592422

[23] Lee, J. W., & Choi, K. W. (2026). Two-tier PEFT adapter switching for environment- and SNR-adaptive channel estimation. IEEE Open Journal of the Communications Society, 7, 6024–6046. https://doi.org/10.1109/ojcoms.2026.3697928

[24] Das, A., Jesbin, F., & Chockalingam, A. (2026). A Gaussian-sinc pulse shaping filter for Zak-OTFS. IEEE Transactions on Vehicular Technology, 75(3), 4307–4322. https://doi.org/10.1109/tvt.2025.3609781

[25] Zhou, M., Chen, F., Xia, M., Zhang, X., & Yu, H. (2025). Iterative channel estimation for multi-user OTFS uplink systems with superimposed full pilots. IEEE Transactions on Vehicular Technology, 74(3), 4485–4497. https://doi.org/10.1109/tvt.2024.3497176

[26] Ying, D., & Ye, F. (2025). Deep learning supported path prediction and channel estimation for MIMO-OTFS system with high delay resolution. IEEE Transactions on Vehicular Technology, 74(3), 3584–3597. https://doi.org/10.1109/tvt.2024.3493921

[27] Mehrotra, A., Srivastava, S., Shanmughanadha Reddy, N., Jagannatham, A., & Hanzo, L. (2025). Sparse channel estimation for MIMO OTFS/OTSM systems using finite-resolution ADCs. IEEE Transactions on Communications, 73(6), 3971–3987. https://doi.org/10.1109/tcomm.2024.3502682

[28] Zhang, X., Liu, C., Yuan, W., Zhang, J. A., & Ng, D. W. K. (2024). Sparse prior-guided deep learning for OTFS channel estimation. IEEE Transactions on Vehicular Technology, 73(12), 19913–19918. https://doi.org/10.1109/tvt.2024.3450012

[29] Li, X., Fan, P., Wang, Q., & Liu, Z. (2025). Grid evolution for doubly fractional channel estimation in OTFS systems. IEEE Transactions on Vehicular Technology, 74(2), 3486–3490. https://doi.org/10.1109/tvt.2024.3472695

[30] Kong, L., Liu, X., Zhang, X., Xiong, J., Zhao, H., & Wei, J. (2025). Representation-based continual learning for channel estimation in dynamic wireless environments. IEEE Transactions on Wireless Communications, 24(8), 6382–6396. https://doi.org/10.1109/twc.2025.3552836

[31] Ma, M., Dai, J., & Jiang, X.-Q. (2025). Fast burst-sparsity learning approach for massive MIMO-OTFS channel estimation. IEEE Wireless Communications Letters, 14(5), 1546–1550. https://doi.org/10.1109/lwc.2025.3548685

[32] Ito, K., Takahashi, T., Furuta, K., Ibi, S., & Thadeu Freitas de Abreu, G. (2024). Joint channel and data estimation via parametric bilinear inference for OTFS demodulation. IEEE Open Journal of the Communications Society, 5, 6090–6105. https://doi.org/10.1109/ojcoms.2024.3463202

[33] Yue, Y., Shi, J., Li, Z., Hu, J., & Tie, Z. (2024). Model-driven deep learning assisted detector for OTFS with channel estimation error. IEEE Communications Letters, 28(4), 842–846. https://doi.org/10.1109/lcomm.2024.3362970