Signal Recovery for Simultaneous Time and Space Erasures

Noah Ernst

Alex Iosevich

April, 2026

Abstract

We formulate a precise statement for the problem of signal recovery when outputs in both time and frequency are erased. From this formulation we find a slight improvement on the deterministic recovery conditions from the Fourier Uncertainty Principle, and characterize the structural reasons for this improvement. We go through further improvements in signal recovery under added conditions. Finally, we prove that our new structural information has been used completely, and no further improvements remain.

Contents

1 Introduction

1.1 Discrete Fourier Analysis

1.1.1 Characters

We define the *primitive character* $\chi : \mathbb{Z}_N \to \mathbb{C}$ by

$$\chi(t) = e^{2\pi i t/N}.$$

For each $m \in \mathbb{Z}_N$, we define the associated character $\chi_m : \mathbb{Z}_N \to \mathbb{C}$ by

$$\chi_m(x) = \chi(mx) = e^{2\pi i m x/N}.$$

These functions are called the *characters* of $\mathbb{Z}_N$. They satisfy the orthogonality relation

$$\sum_{x \in \mathbb{Z}_N} \chi_m(x)\overline{\chi_{m'}(x)} = \begin{cases} N, & m = m', \\ 0, & m \neq m'. \end{cases}$$

In particular, the collection

$$\left\{ \frac{1}{\sqrt{N}} \chi_m \right\}_{m \in \mathbb{Z}_N}$$

forms an orthonormal set with respect to the standard inner product on functions $f : \mathbb{Z}_N \to \mathbb{C}$.

1.1.2 Discrete Fourier Transform

We equip the space of functions $f : \mathbb{Z}_N \to \mathbb{C}$ with the inner product

$$\langle f, g \rangle = \sum_{x \in \mathbb{Z}_N} f(x)\overline{g(x)}.$$

Since the normalized characters form an orthonormal basis, any function f admits the expansion

$$f = \sum_{m \in \mathbb{Z}_N} \left\langle f, \frac{1}{\sqrt{N}} \chi_m \right\rangle \frac{1}{\sqrt{N}} \chi_m.$$

This motivates the definition of the *Fourier transform* of f as

$$\widehat{f}(m) = \left\langle f, \frac{1}{\sqrt{N}} \chi_m \right\rangle.$$

Expanding the inner product, we obtain

$$\widehat{f}(m) = \frac{1}{\sqrt{N}} \sum_{x \in \mathbb{Z}_N} f(x)\overline{\chi_m(x)} = \frac{1}{\sqrt{N}} \sum_{x \in \mathbb{Z}_N} f(x)\chi(-mx).$$

1.1.3 Fourier Inversion

By the orthonormal basis expansion,

$$f = \sum_{m \in \mathbb{Z}_N} \widehat{f}(m) \frac{1}{\sqrt{N}} \chi_m.$$

Evaluating at $x \in \mathbb{Z}_N$ gives

$$f(x) = \frac{1}{\sqrt{N}} \sum_{m \in \mathbb{Z}_N} \widehat{f}(m) \chi_m(x) = \frac{1}{\sqrt{N}} \sum_{m \in \mathbb{Z}_N} \widehat{f}(m) e^{2\pi i m x/N}.$$

This is the *Fourier inversion formula*.

1.1.4 Plancherel Identity

Let $f, g : \mathbb{Z}_N \to \mathbb{C}$. Then

$$\langle \widehat{f}, \widehat{g} \rangle = \sum_{m \in \mathbb{Z}_N} \widehat{f}(m) \overline{\widehat{g}(m)}.$$

Substituting the definition of the Fourier transform gives

$$\langle \widehat{f}, \widehat{g} \rangle = \frac{1}{N} \sum_{x,y \in \mathbb{Z}_N} f(x) \overline{g(y)} \sum_{m \in \mathbb{Z}_N} \chi(m(y-x)).$$

By orthogonality of characters,

$$\sum_{m \in \mathbb{Z}_N} \chi(m(y-x)) = \begin{cases} N, & x = y, \\ 0, & x \neq y. \end{cases}$$

Therefore,

$$\langle \widehat{f}, \widehat{g} \rangle = \sum_{x \in \mathbb{Z}_N} f(x) \overline{g(x)} = \langle f, g \rangle.$$

This proves the *Plancherel identity*:

$$\langle f, g \rangle = \langle \widehat{f}, \widehat{g} \rangle.$$

1.1.5 Fourier Uncertainty Principle

Definition 1.1. *For $f : \mathbb{Z}_N \to \mathbb{C}$, the support of f is*

$$\operatorname{supp}(f) = \{n \in \mathbb{Z}_N : f(n) \neq 0\}.$$

The Fourier uncertainty principle states that a function and its Fourier transform cannot be simultaneously localized. In one dimension, for a nonzero function $f : \mathbb{Z}_N \to \mathbb{C}$,

$$|\operatorname{supp}(f)| \cdot |\operatorname{supp}(\widehat{f})| \geq N.$$

Theorem 1.1 (Discrete uncertainty principle)**.** *Let $f : \mathbb{Z}_N \to \mathbb{C}$ be a nonzero function supported on $E \subset \mathbb{Z}_N$, and suppose its Fourier transform $\widehat{f}$ is supported on $S \subset \mathbb{Z}_N$. Then*

$$|E| \cdot |S| \geq N.$$

Proof. Since $\widehat{f}$ is supported on S, the inverse Fourier transform gives

$$f(x) = \frac{1}{N} \sum_{m \in S} \widehat{f}(m)\, \chi(xm).$$

Using the definition of the Fourier transform,

$$\widehat{f}(m) = \sum_{x' \in E} f(x')\, \chi(-x'm),$$

because f vanishes outside E.
Substituting this into the inverse formula yields

$$f(x) = \frac{1}{N} \sum_{m \in S} \sum_{x' \in E} f(x')\, \chi(xm)\chi(-x'm).$$

Hence

$$|f(x)| \leq \frac{1}{N} \sum_{m \in S} \sum_{x' \in E} |f(x')|.$$

Since the sums are finite,

$$|f(x)| \leq \frac{|S|}{N} \sum_{x' \in E} |f(x')|.$$

Summing over all $x \in E$ gives

$$\sum_{x \in E} |f(x)| \leq \frac{|E||S|}{N} \sum_{x' \in E} |f(x')|.$$

Because $f \neq 0$, the quantity

$$\sum_{x \in E} |f(x)|$$

is strictly positive, and therefore we may divide both sides by it to obtain

$$1 \leq \frac{|E||S|}{N}.$$

Thus

$$|E| \cdot |S| \geq N,$$

as claimed. □

2 One-Sided Signal Recovery

2.1 ℓ_1 Minimization

Let $f : \mathbb{Z}_N \to \mathbb{C}$ be a one-dimensional signal. Suppose that we observe its Fourier transform $\widehat{f}$ except on a set $S \subset \mathbb{Z}_N$ of missing frequencies:

$$\widehat{f}(m) = N^{-1/2} \sum_{x \in \mathbb{Z}_N} f(x)\chi(-xm), \quad m \notin S.$$

We aim to recover f by solving

$$g = \operatorname*{arg\,min}_{u : \widehat{u}(m) = \widehat{f}(m) \text{ for } m \notin S} \|u\|_1.$$

Let

$$h = f - g.$$

By construction, $\widehat{h}(m) = 0$ for all $m \notin S$.
Let $E = \operatorname{supp}(f)$. Then

$$\|g\|_1 = \|f - h\|_1 = \|f - h\|_{L^1(E)} + \|h\|_{L^1(E^c)}.$$

By the triangle inequality,

$$\|g\|_1 \geq \|f\|_1 + \left(\|h\|_{L^1(E^c)} - \|h\|_{L^1(E)}\right).$$

Using Fourier inversion, for $x \in E$,

$$h(x) = N^{-1/2} \sum_{m \in S} \widehat{h}(m)\chi(xm).$$

Thus

$$|h(x)| \leq N^{-1/2}|S| \max_{m \in S} |\widehat{h}(m)|.$$

Since

$$|\widehat{h}(m)| \leq N^{-1/2}\|h\|_1,$$

we have

$$|h(x)| \leq N^{-1}|S|\|h\|_1.$$

Therefore,

$$\|h\|_{L^1(E)} \leq N^{-1}|E||S|\|h\|_1.$$

If

$$|E||S| < \frac{N}{2},$$

then

$$\|h\|_{L^1(E)} < \frac{1}{2}\|h\|_1,$$

so

$$\|h\|_{L^1(E^c)} > \|h\|_{L^1(E)}.$$

Thus

$$\|g\|_1 > \|f\|_1,$$

which contradicts the fact that g is the ℓ_1 minimizer unless $h \equiv 0$. Hence $g = f$.
Therefore, under the condition

$$|E||S| < \frac{N}{2},$$

ℓ_1 minimization exactly recovers f.

2.2 Donoho–Stark Recovery

The ℓ_1 minimization method provides both a proof and a recovery system for signals. However, there is an easier method of proving unique recovery using the Fourier uncertainty principle. [1]
Suppose that f is supported in $E \subset \mathbb{Z}_N$ and that $\widehat{f}$ is observed except on a set $S \subset \mathbb{Z}_N$. Let g be another candidate signal such that

$$\widehat{g}(m) = \widehat{f}(m), \quad m \notin S.$$

Set

$$h = f - g.$$

Then $\widehat{h}$ is supported in S. If g is also supported on a set of size at most $|E|$, then h is supported on a set of size at most $2|E|$. Hence, if $h \neq 0$, the uncertainty principle gives

$$2|E||S| \geq N.$$

Therefore, if

$$|E||S| < \frac{N}{2},$$

then $h = 0$, so $f = g$.

With this one-sided result established we progress to the main problem addressed in this thesis.

3 Two-Sided Signal Recovery

3.1 Problem formulation

Let $X \subset \mathbb{Z}_N$ be a fixed set of frequencies. Consider signals

$$f : \mathbb{Z}_N \to \mathbb{C}$$

satisfying

$$\text{supp}(\widehat{f}) \subset X.$$

Let $M \subset \mathbb{Z}_N$ be a set of erased time samples. $f(n)$ is known for all $n \in M^c$, and the recovery problem is to determine whether these observed values uniquely determine f.

Without additional assumptions this is not possible in general. The key hypothesis is that the Fourier transform of f is supported on a sufficiently small set. This setting differs from the classical Donoho–Stark problem in an essential way. In the classical formulation one assumes a signal sparse in time and partially observes its Fourier transform. Here the signal is partially observed in time, and sparsity appears only after the decomposition

$$f = g + h,$$

where

$$g = f \cdot \mathbf{1}_{M^c}, \qquad \operatorname{supp}(h) \subset M.$$

3.2 Deterministic recovery bound

Theorem 3.1 (Two-sided recovery bound)**.** *Let $M, X \subset \mathbb{Z}_N$. Suppose*

$$\operatorname{supp}(\widehat{f}) \subset X.$$

If

$$|M||X| < N,$$

then f is uniquely determined by its values on M^c.

Proof. Suppose f_1 and f_2 satisfy

$$\operatorname{supp}(\widehat{f_1}), \operatorname{supp}(\widehat{f_2}) \subset X,$$

and

$$f_1(n) = f_2(n) \quad \text{for all } n \in M^c.$$

Let

$$h = f_1 - f_2.$$

Then

$$\operatorname{supp}(h) \subset M, \qquad \operatorname{supp}(\widehat{h}) \subset X.$$

If $h \neq 0$, the discrete uncertainty principle gives

$$|\operatorname{supp}(h)|\,|\operatorname{supp}(\widehat{h})| \geq N.$$

But

$$|\operatorname{supp}(h)| \leq |M|, \qquad |\operatorname{supp}(\widehat{h})| \leq |X|,$$

so

$$|\operatorname{supp}(h)|\,|\operatorname{supp}(\widehat{h})| \leq |M||X| < N,$$

a contradiction. Therefore $h = 0$. □

Discussion. The improvement over the classical one-sided Donoho–Stark bound comes from a sharper estimate on the support of the error signal. In the one-sided setting the difference between two candidate reconstructions may lie anywhere in the union of their supports, producing a factor of two in the classical bound. In the two-sided setting, however, the signals already agree outside the erased time set M, so the error signal is supported entirely inside M.

Applying the uncertainty principle to this smaller support yields the improved condition

$$|M||X| < N.$$

3.3 Sharpness and saturation of the uncertainty principle

In order to formulate more effective recovery conditions, it is instructive to investigate the structures under which our recovery bound is sharp, i.e. recovery fails at $|M||X| < N$.

Theorem 3.2 (General conditions for failure)**.** *Recovery fails for a pair (M, X) if and only if there exists a nonzero function h supported on the erased set whose Fourier transform is supported on the allowed frequency set, that is,*

$$\text{supp}(h) \subseteq M, \qquad \text{supp}(\widehat{h}) \subseteq X.$$

In particular, the uncertainty principle implies that this can occur only when

$$|M||X| \geq N.$$

Proof. Suppose recovery fails. Then there exist two distinct signals f_1 and f_2 that agree on M^c and both have Fourier support contained in X. Define

$$h = f_1 - f_2.$$

Then $h \neq 0$ and

$$\text{supp}(h) \subseteq M, \qquad \text{supp}(\widehat{h}) \subseteq X.$$

Conversely, suppose there exists a nonzero function h satisfying

$$\text{supp}(h) \subseteq M, \qquad \text{supp}(\widehat{h}) \subseteq X.$$

If f is any signal whose Fourier transform is supported in X, then f and $f + h$ agree on M^c and both satisfy the same Fourier support constraint. Thus the observed data on M^c cannot distinguish between these signals, and recovery fails.
Applying the discrete uncertainty principle to h gives

$$|\,\text{supp}(h)|\,|\,\text{supp}(\widehat{h})| \geq N,$$

so failure of recovery can occur only when

$$|M||X| \geq N.$$

□

Discussion. This result identifies the precise mechanism behind failure of recovery. Non-uniqueness occurs exactly when there exists a nonzero signal that is simultaneously sparse in both the erased time indices and the allowed frequency set. In particular, the deterministic recovery condition

$$mx < N$$

follows directly from the discrete Fourier uncertainty principle.

Theorem 3.3 (Structured subgroup overlap obstruction)**.** *Suppose $m \mid N$. Then there exist sets $M, X \subset \mathbb{Z}_N$ with*

$$|M| = m, \qquad |X| = \frac{N}{m}, \qquad |M||X| = N,$$

for which recovery from the samples on M^c fails for signals whose Fourier transform is supported in X. More precisely, this failure occurs when M is a subgroup of $\mathbb{Z}_N$ of size m, and X is its annihilator subgroup.

Proof. Let $H \subset \mathbb{Z}_N$ be a subgroup with $|H| = m$, and define

$$h = 1_H.$$

Then

$$\operatorname{supp}(h) = H, \qquad |\operatorname{supp}(h)| = m.$$

Consider the annihilator subgroup

$$H^\perp = \{k \in \mathbb{Z}_N : \chi(kn) = 1 \text{ for every } n \in H\}.$$

A standard computation using orthogonality of characters shows that

$$\widehat{1_H}(k) = c\, 1_{H^\perp}(k)$$

for some nonzero constant c, and therefore

$$\operatorname{supp}(\widehat{h}) = H^\perp.$$

Moreover,

$$|H^\perp| = \frac{N}{m}.$$

Now set

$$M = H, \qquad X = H^\perp.$$

Then

$$|M| = m, \qquad |X| = \frac{N}{m}, \qquad |M||X| = N.$$

The function h is nonzero and satisfies

$$\operatorname{supp}(h) \subset M, \qquad \operatorname{supp}(\widehat{h}) \subset X.$$

Hence the two signals

$$f_1 = 0, \qquad f_2 = h$$

both have Fourier support contained in X and agree on all observed samples M^c, since

$$h(n) = 0 \quad \text{for every } n \in M^c.$$

However $f_1 \neq f_2$, so the samples on M^c do not uniquely determine the signal. □

Discussion. The obstruction in Theorem 3.3 arises from a highly structured alignment between the missing time indices and the allowed Fourier support. Specifically, the counterexample is constructed by choosing a subgroup

$$H \subset \mathbb{Z}_N$$

of size m, and considering its indicator function

$$h = \mathbf{1}_H.$$

This function is supported exactly on H, which is a regularly spaced set of indices in $\mathbb{Z}_N$. The key feature of subgroup indicator functions is that their Fourier transforms are also supported on a subgroup, namely the annihilator subgroup

$$H^\perp = \{k \in \mathbb{Z}_N : \chi(kn) = 1 \text{ for every } n \in H\}.$$

This set consists precisely of those frequencies that are compatible with the periodic spacing of H. In particular,

$$|\operatorname{supp}(h)| = |H| = m, \qquad |\operatorname{supp}(\widehat{h})| = |H^\perp| = \frac{N}{m},$$

so

$$|\operatorname{supp}(h)|\,|\operatorname{supp}(\widehat{h})| = N.$$

Thus subgroup indicator functions exactly saturate the discrete Fourier uncertainty principle. This illustrates the effectiveness of Theorem 3.2, showing a concrete example of a structure where recovery fails at $mx = N$

It is important to note for future improvements that this failure of recovery occurs only because the sets M and X are chosen in a very special way. Later improvements to the recovery bound will rely on excluding this type of structured alignment between the missing time indices and the frequency support through strategies such as randomly generated supports.

Example 3.1 (Exact support containment is necessary for failure of recovery in Theorem 3.2)**.** We illustrate why the condition

$$\operatorname{supp}(h) \subseteq M, \qquad \operatorname{supp}(\widehat{h}) \subseteq X$$

must hold exactly in order for recovery to fail.

Let $N = 5$, and work in $\mathbb{Z}_5$. Suppose the erased time samples are

$$M = \{1, 2\},$$

so the observed samples are

$$M^c = \{0, 3, 4\}.$$

Let the allowed frequency set be

$$X = \{0, 1\}.$$

Assume that signals satisfy the Fourier support condition

$$\operatorname{supp}(\widehat{f}) \subseteq X.$$

Now consider a function h with support

$$\operatorname{supp}(h) = \{1, 2, 3\}.$$

Thus h is supported mostly inside M, but not entirely, since

$$3 \in M^c.$$

Suppose that two candidate reconstructions satisfy

$$f_2 = f_1 + h.$$

Then at the observed index $3 \in M^c$ we have

$$f_2(3) = f_1(3) + h(3) \neq f_1(3),$$

so the two signals disagree on observed data. Therefore they are distinguishable from the available samples, and recovery remains unique.

Discussion. This example shows that failure of recovery does not occur when a signal is *almost* supported inside the erased time indices. If even one point of $\operatorname{supp}(h)$ lies outside M, then the observed samples detect the difference between two candidate reconstructions. An analogous statement holds on the Fourier side. If

$$\operatorname{supp}(\widehat{h}) \not\subseteq X,$$

then there exists a frequency $k \notin X$ such that

$$\widehat{h}(k) \neq 0,$$

and therefore one of the candidate signals fails to satisfy the admissible Fourier support condition. Hence it cannot serve as an alternative reconstruction.
Thus recovery fails precisely when there exists a nonzero function h such that

$$\operatorname{supp}(h) \subseteq M, \qquad \operatorname{supp}(\widehat{h}) \subseteq X.$$

In this case both f and $f + h$ satisfy the Fourier support constraint and agree on all observed samples, making them indistinguishable from the available data. This explains why exact support containment characterizes the mechanism of non-uniqueness.

Example 3.2 (Exact containment is impossible when $mx < N$)**.** Let $N = 5$, and work in $\mathbb{Z}_5$. Suppose

$$M = \{1\}, \qquad X = \{0, 1, 2, 3\}.$$

Then

$$|M||X| = 1 \cdot 4 = 4 < 5 = N.$$

We claim that no nonzero function $h : \mathbb{Z}_5 \to \mathbb{C}$ can satisfy

$$\operatorname{supp}(h) \subseteq M, \qquad \operatorname{supp}(\widehat{h}) \subseteq X.$$

Indeed, if $\operatorname{supp}(h) \subseteq M = \{1\}$, then h has the form

$$h(n) = c\,\mathbf{1}_{\{1\}}(n)$$

for some constant $c \in \mathbb{C}$. Its Fourier transform is

$$\widehat{h}(k) = \frac{1}{\sqrt{5}} \sum_{n \in \mathbb{Z}_5} h(n) e^{-2\pi i n k/5} = \frac{c}{\sqrt{5}} e^{-2\pi i k/5}.$$

If $c \neq 0$, then

$$\widehat{h}(k) \neq 0$$

for every $k \in \mathbb{Z}_5$. Hence

$$\operatorname{supp}(\widehat{h}) = \mathbb{Z}_5,$$

which is not contained in

$$X = \{0, 1, 2, 3\}.$$

Therefore the only function satisfying both

$$\operatorname{supp}(h) \subseteq M, \qquad \operatorname{supp}(\widehat{h}) \subseteq X$$

is $h \equiv 0$. By the failure of Theorem 3.2, recovery is unique for this pair (M, X).

Discussion. This example shows concretely how the condition $mx < N$ prevents failure. Although X contains most of the frequencies, the time support M is so small that any nonzero function supported on M must have Fourier transform supported on all of $\mathbb{Z}_5$. Thus it cannot also be supported inside X.

In other words, exact containment in both domains is too restrictive when

$$|M||X| < N.$$

The uncertainty principle rules out the existence of a nonzero obstruction h, and therefore recovery is unique. With this clearly established, we also find an important generalization, which will apply nicely to Theorem 3.3

Theorem 3.4 (Translation and modulation)**.** *If recovery fails for a pair (M, X) satisfying $mx = N$, then it also fails for any translate of M and any modulation of the corresponding signal. Thus the subgroup alignment issues persist across entire symmetry classes of supports.*

Proof. Let $h = \mathbf{1}_H$ be the subgroup example above. For any $a, b \in \mathbb{Z}_N$, define

$$h_{a,b}(n) = \chi(bn) h(n - a).$$

Translation preserves the size of the time support, and modulation preserves the size of the frequency support. Hence

$$|\operatorname{supp}(h_{a,b})| = m, \qquad |\operatorname{supp}(\widehat{h_{a,b}})| = \frac{N}{m}.$$

Thus each $h_{a,b}$ also saturates the uncertainty principle and produces non-uniqueness at the endpoint $mx = N$.

□

Discussion. The specific subgroup case which does not allow for recovery of our signal generates an entire orbit of such cases under translation and modulation, which are the natural symmetries of the discrete Fourier transform. Consequently, failure of recovery at the endpoint $mx = N$ cannot be attributed to the particular location of the missing indices, but rather to the presence of an underlying subgroup-type structure that remains invariant under these symmetries. This observation clarifies the mechanism behind the sharpness of the deterministic recovery bound. The endpoint obstruction cannot be removed by simply repositioning the missing samples or frequency constraints. This strengthens our later use of randomized supports, as a "simpler" seeming fix of moving around our supports would fail under Theorem 3.4

4 Additive Energy

4.1 Refinement of the uncertainty principle via additive structure

The deterministic recovery condition $|M||X| < N$ arises from the discrete Fourier uncertainty principle and is sharp in general. As shown in Section 3.3, equality is attained by highly structured configurations, specifically subgroup–annihilator pairs, where a nonzero function is simultaneously localized in time and frequency. However, this bound depends only on the cardinalities of the supports and does not distinguish between different types of frequency sets. In particular, it treats a random set and a highly structured subgroup of the same size as equivalent, even though their Fourier behavior is very different. To formalize this, we introduce *additive energy*, a quantitative measure of additive structure. This allows us to strengthen the uncertainty principle in situations where the frequency support lacks strong internal structure.

Definition 4.1. *Let $X \subset \mathbb{Z}_N$. The* additive energy *of X is defined by*

$$\Lambda(X) = \left|\{(x_1, x_2, x_3, x_4) \in X^4 : x_1 + x_2 = x_3 + x_4\}\right|.$$

This counts the number of additive relations inside X.

Discussion. If X has many additive relations (e.g. a subgroup or arithmetic progression), then $\Lambda(X)$ is large, on the order of $|X|^3$. If X behaves like a random set, then additive coincidences are rare, and

$$\Lambda(X) \sim |X|^2.$$

Thus additive energy measures how far a set is from the highly structured subgroup obstruction identified earlier [2].

A central identity linking additive energy to Fourier analysis is

$$\Lambda(X) = N^3 \sum_{m \in \mathbb{Z}_N} |\widehat{\mathbf{1}_X}(m)|^4,$$

so large additive energy corresponds to concentration of Fourier mass. This is the mechanism by which additive structure influences uncertainty.

4.2 Additive energy uncertainty principle

We now state the refined uncertainty principle.

Theorem 4.1 (Additive uncertainty principle [2]). *Let $h \neq 0$ satisfy*

$$\operatorname{supp}(h) \subset E, \qquad \operatorname{supp}(\widehat{h}) \subset X.$$

Then

$$N \leq |E| \cdot \Lambda^{1/3}(X).$$

This improves the classical bound

$$N \leq |E||X|$$

by replacing $|X|$ with the smaller quantity $\Lambda^{1/3}(X)$.

4.3 Recovery theorem

Theorem 4.2 (Additive energy recovery condition). *Let $M, X \subset \mathbb{Z}_N$, and suppose*

$$\operatorname{supp}(\widehat{f}) \subset X.$$

If

$$|M| \cdot \Lambda^{1/3}(X) < N,$$

then f is uniquely determined by its values on M^c.

Proof. Suppose f_1 and f_2 agree on M^c and both satisfy

$$\operatorname{supp}(\widehat{f_1}), \operatorname{supp}(\widehat{f_2}) \subset X.$$

Let

$$h = f_1 - f_2.$$

Then

$$\operatorname{supp}(h) \subset M, \qquad \operatorname{supp}(\widehat{h}) \subset X.$$

If $h \neq 0$, the additive uncertainty principle gives

$$N \leq |\operatorname{supp}(h)| \cdot \Lambda^{1/3}(X).$$

Since $|\operatorname{supp}(h)| \leq |M|$, we obtain

$$N \leq |M| \cdot \Lambda^{1/3}(X),$$

which contradicts the hypothesis. Therefore $h = 0$, and recovery is unique. □

Discussion. Since

$$\Lambda(X) \leq |X|^3,$$

we always have

$$\Lambda^{1/3}(X) \leq |X|.$$

Thus the additive energy condition

$$|M|\Lambda^{1/3}(X) < N$$

is always at least as strong as

$$|M||X| < N.$$

If X is maximally structured (e.g. a subgroup), then

$$\Lambda(X) \sim |X|^3,$$

so the bound reduces to

$$|M||X| < N.$$

This recovers exactly the subgroup obstruction from Section 3.3. Additive energy does not remove the sharp examples, but it detects them, which allows us to make a strictly better bound (in cases that are not sharp) without fear of missing the sharp cases, as is explicitly seen below.

On the other hand, if X has little additive structure, then additive coincidences are rare. For instance, for a random subset X of size $|X| = k$, one expects that a relation

$$x_1 + x_2 = x_3 + x_4$$

occurs only when the pairs (x_1, x_2) and (x_3, x_4) are essentially the same, leading to

$$\Lambda(X) \sim |X|^2.$$

and the condition becomes

$$|M| \cdot |X|^{2/3} < N,$$

which is strictly stronger than the classical bound.

Intuitively, since the classical uncertainty principle is purely *cardinality-based* it fails to distinguish between highly structured sets (subgroups which make recovery difficult) and unstructured sets (random supports which make recovery easy). Additive energy measures how much additive structure exists inside X, and therefore how close we are to worst-case recovery scenarios. Based on this information, the bound can be made more sophisticated rather than accommodating the possibility of high structure in all cases.

5 Improvements in special cases

Under the condition that N is prime, we can find a stronger version of the Fourier uncertainty principle, especially in cases when N is large.

Theorem 5.1 (Additive uncertainty principle in $\mathbb{Z}_N$ for N prime)**.** *Assume N is prime. If $h \not\equiv 0$, then*

$$|\operatorname{supp}(h)| + |\operatorname{supp}(\widehat{h})| \geq N + 1.$$

Proof. Let

$$k = |\operatorname{supp}(h)|.$$

Then h vanishes at $N - k$ points of $\mathbb{Z}_N$.
The Fourier transform $\widehat{h}$ is a trigonometric polynomial of degree at most $N-1$. If $\widehat{h}$ vanishes on more than

$$N - k - 1$$

points, then it must vanish identically. This would imply

$$h \equiv 0,$$

contradicting the assumption that $h \not\equiv 0$.
Therefore,

$$|\operatorname{supp}(\widehat{h})| \geq N + 1 - k.$$

Rearranging gives

$$|\operatorname{supp}(h)| + |\operatorname{supp}(\widehat{h})| \geq N + 1,$$

□

Theorem 5.2 (Prime modulus recovery condition)**.** *Assume N is prime. If*

$$|M| + |X| \leq N,$$

then recovery holds

Proof. Suppose recovery fails. Then there exists a nonzero function h such that

$$\operatorname{supp}(h) \subseteq M, \qquad \operatorname{supp}(\widehat{h}) \subseteq X.$$

Hence

$$|\operatorname{supp}(h)| \leq |M|, \qquad |\operatorname{supp}(\widehat{h})| \leq |X|.$$

Applying the previous theorem gives

$$|M| + |X| \geq N + 1,$$

which contradicts the assumption

$$|M| + |X| \leq N.$$

Therefore recovery must be unique. □

Discussion. Intuitively, this improvement comes from the fact that when N is prime, the group $\mathbb{Z}_N$ has no nontrivial subgroups. The counterexamples that cause failure at the endpoint in the general case rely on this kind of subgroup structure, so they cannot occur here. As a result, the usual product condition $|M||X| < N$ can be strengthened in the prime setting.

6 Linear Algebra Formulation

Let

$$V_X = \{f : \mathrm{supp}(\widehat{f}) \subset X\}.$$

This space represents the collection of all signals compatible with the frequency-side constraint in the recovery problem.

Theorem 6.1. *Recovery from the samples on M^c is unique if and only if the restriction map*

$$R : V_X \to \mathbb{C}^{M^c}, \qquad R(f) = f|_{M^c}$$

is injective.

Proof. Injectivity of R means that if $f \in V_X$ satisfies

$$f|_{M^c} = 0,$$

then $f = 0$. The condition $f|_{M^c} = 0$ is equivalent to

$$\mathrm{supp}(f) \subset M.$$

Since $f \in V_X$, we also have

$$\mathrm{supp}(\widehat{f}) \subset X.$$

Thus injectivity holds if and only if there is no nonzero function satisfying

$$\mathrm{supp}(f) \subset M, \qquad \mathrm{supp}(\widehat{f}) \subset X.$$

This is exactly the condition for uniqueness of recovery. □

Discussion. Rather than asking directly whether a signal f is uniquely determined by its values on M^c, we instead ask whether the restriction operator

$$R : V_X \to \mathbb{C}^{M^c}, \qquad R(f) = f|_{M^c},$$

is injective. Recovery is equivalent to the statement that no nonzero signal in V_X can vanish on the observed indices. Thus the recovery problem reduces to determining whether the intersection

$$V_X \cap V_M$$

is trivial, where

$$V_M = \{f : \mathrm{supp}(f) \subseteq M\}.$$

Again, recovery fails exactly when there exists a nonzero signal that is simultaneously localized inside the erased time indices and inside the allowed frequency set. This perspective clarifies the role of the uncertainty principle in the previous section. The condition $|M||X| < N$ ensures that such a signal cannot exist, and therefore guarantees that the restriction map R is injective. In this way, the uncertainty principle can be interpreted as providing a sufficient condition for injectivity of the restriction operator.

This operator viewpoint is useful because it allows recovery to be studied using geometric and probabilistic methods. In particular, we use this framework to show that even when the deterministic condition $|M||X| < N$ fails, the restriction map may still be injective for typical choices of sets M and X.

7 Random sampling

Our recovery condition is sharp in general due to subgroup-type alignments between the erased time indices and admissible frequency supports [Section 3.3]. However, these obstructions are highly structured and therefore unstable under generic perturbations of the sampling sets. This suggests that stronger recovery guarantees should hold when the supports are chosen randomly. We now show that this intuition can be made rigorous. In particular, randomly selected sampling sets behave in a manner analogous to orthogonal systems when restricted to sparse vectors, which prevents the existence of simultaneous time–frequency localization inside the pair (M, X).

7.1 Random Fourier restriction operators

Recall from Section 6 that recovery from the samples on M^c is unique if and only if the restriction operator

$$R : V_X \to \mathbb{C}^{M^c}, \qquad R(f) = f|_{M^c}$$

is injective. Thus, recovery fails if and only if there exists $h \neq 0$ such that

$$\operatorname{supp}(h) \subseteq M, \qquad \operatorname{supp}(\widehat{h}) \subseteq X.$$

This is equivalent to the existence of a nonzero signal supported on M that lies in the null space of the Fourier transform restricted to the rows indexed by X^c. Random sampling improves recovery by showing with high probability that such null space vectors do not exist for typical choices of the supports.

Theorem 7.1 (Random Fourier recovery)**.** *Let $X \subset \mathbb{Z}_N$ be chosen uniformly at random with $|X| = x$, and suppose*

$$\operatorname{supp}(\widehat{f}) \subseteq X.$$

Then with high probability over the choice of X, every signal supported on a set $M \subset \mathbb{Z}_N$ satisfying

$$|M| \lesssim \frac{x}{(\log N)^6}$$

is uniquely determined by its values on M^c.

Proof. Recovery fails if and only if there exists a nonzero function

$$h \neq 0$$

such that

$$\operatorname{supp}(h) \subseteq M, \qquad \operatorname{supp}(\widehat{h}) \subseteq X.$$

Equivalently,

$$\widehat{h}|_{X^c} = 0.$$

Let A denote the restriction of the discrete Fourier transform to the rows indexed by X^c. Then the condition above implies

$$Ah = 0.$$

Thus recovery fails precisely when the restricted Fourier matrix has a nontrivial null vector supported on M. A result from [3] shows that if the rows of the discrete Fourier transform are selected uniformly at random, then with high probability the resulting restricted operator A approximately preserves the ℓ^2 norm of all vectors supported on sufficiently small sets. More precisely, there exist constants $c, C > 0$ such that if

$$|M| \leq C\frac{x}{(\log N)^6},$$

then every vector v supported on M satisfies

$$c\|v\|_2 \leq \|Av\|_2 \leq C\|v\|_2.$$

In particular, such a vector cannot lie in the null space of A unless $v = 0$. Applying this to h shows that $h = 0$, and therefore recovery is unique.

□

Discussion. In contrast to our worst case scenarios, we see that when the frequency support is chosen randomly, the restricted Fourier operator behaves almost like an orthogonal system on sparse vectors. As a result, simultaneous localization inside (M, X) becomes extremely unlikely due to a previously known result, and recovery holds for substantially larger values of $|M|$ than predicted by the deterministic bound.

7.2 Improvement specific to the two-sided setting

It is important to re-emphasize the differences of our current problem from the classical one-sided formulation considered in Donoho–Stark type results. As noted before, in the one-sided setting, two candidate signals may differ on the union of their supports. In contrast, in the two-sided setting considered here, the difference between two candidate signals is already supported entirely inside the erased set M. Consequently, random sampling only needs to rule out the existence of functions satisfying

$$\text{supp}(h) \subseteq M, \qquad \text{supp}(\widehat{h}) \subseteq X,$$

rather than controlling a larger support union as in the one-sided case. This structural advantage leads to sharper recovery guarantees and makes probabilistic injectivity of the restriction operator easier to obtain.

8 Concentration Inequalities

9 Exhaustion of the structural information in the two-sided model

We now show that the two-sided observation model introduces exactly one additional structural constraint compared with the classical one-sided setting, namely the confinement of the error signal to the erased time indices. No further universal restriction on admissible differences between candidate reconstructions can be deduced from the model, therefore the only extractable structural improvement is the one discussed following Theorem 3.1

Theorem 9.1 (Maximality of the two-sided support constraint)**.** *Let $M, X \subset \mathbb{Z}_N$. Suppose that f_1, f_2 satisfy*

$$\operatorname{supp}(\widehat{f_1}), \operatorname{supp}(\widehat{f_2}) \subset X$$

and

$$f_1 = f_2 \quad \text{on } M^c.$$

Then their difference

$$h = f_1 - f_2$$

satisfies

$$\operatorname{supp}(h) \subseteq M.$$

Moreover, this containment is the only universal restriction on the support of admissible difference signals imposed by the observation model.

Proof. The agreement of f_1 and f_2 on M^c immediately implies

$$\operatorname{supp}(h) \subseteq M.$$

Conversely, let h be any function satisfying

$$\operatorname{supp}(h) \subseteq M, \qquad \operatorname{supp}(\widehat{h}) \subseteq X.$$

Then the signals

$$f_1 = 0, \qquad f_2 = h$$

are both admissible and agree on M^c.
Thus every function satisfying these support conditions arises as the difference between admissible reconstructions. No further restriction on h follows from the observation model. □

Discussion. This theorem shows that the two-sided formulation contributes exactly one additional piece of structural information beyond the classical setting: the error signal is confined to the erased time indices. The deterministic improvement

$$|M||X| < N$$

obtained earlier follows precisely from this support containment. Once the error signal has been restricted to M, recovery reduces to ruling out functions simultaneously supported in M and Fourier-supported in X.

Since every such function can occur as the difference between admissible signals, the observation model imposes no further universal constraints. Thus the improved recovery condition already exhausts the additional information available from the two-sided setup itself. Any further refinement must therefore rely on additional structure of the sets M or X, rather than on the observation model alone.

Acknowledgments

I want to thank my advisor, Professor Alex Iosevich, as well as all the other math professors I have had during undergrad. I also want to thank my friends from math and from physics, especially Aadit Jain, who wrote a similar thesis.

References

[1] David L. Donoho and Philip B. Stark. Uncertainty principles and signal recovery. *SIAM Journal on Applied Mathematics*, 49(3):906–931, 1989.

[2] K. Aldahleh, A. Iosevich, J. Iosevich, J. Jaimangal, A. Mayeli, and S. Pack. Additive energy, uncertainty principle and signal recovery mechanisms. *arXiv preprint arXiv:2504.14702*, 2025.

[3] Emmanuel Candes, Justin Romberg, and Terence Tao. Stable signal recovery from incomplete and inaccurate measurements. *arXiv preprint*, 2005.