Restriction estimates for hyperbolic paraboloids in higher dimensions via bilinear estimates

Alex Barron

Abstract. Let $\mathbb{H}$ be a $(d-1)$-dimensional hyperbolic paraboloid in $\mathbb{R}^d$ and let Ef be the Fourier extension operator associated to $\mathbb{H}$, with f supported in $B^{d-1}(0,2)$. We prove that $\|Ef\|_{L^p(B(0,R))} \leq C_\varepsilon R^\varepsilon \|f\|_{L^p}$ for all $p \geq 2(d+2)/d$ whenever $d/2 \geq m+1$, where m is the minimum between the number of positive and negative principal curvatures of $\mathbb{H}$. Bilinear restriction estimates for $\mathbb{H}$ proved by S. Lee and Vargas play an important role in our argument.

1. Introduction

In this paper we study estimates for the operator

$$Ef(x,t) = \int_{\mathbb{R}^{d-1}} f(\xi)\, e^{2\pi i(x\cdot\xi + t(\xi_1^2 + \cdots + \xi_{d-m-1}^2 - \xi_{d-m}^2 - \xi_{d-m+1}^2 - \cdots - \xi_{d-1}^2))}\, d\xi,$$

$$\mathrm{supp}(f) \subset B^{d-1}(0,2).$$

This is the extension operator associated to the hyperbolic paraboloid

$$\mathbb{H} = \{\xi \in \mathbb{R}^d : \xi_d = \xi_1^2 + \xi_2^2 + \cdots + \xi_{d-m-1}^2 - \xi_{d-m}^2 - \cdots - \xi_{d-1}^2\}.$$

We let M denote the $(d-1) \times (d-1)$ diagonal matrix with $M_{ii} = 1$ if $i \leq d-1-m$ and $M_{ii} = -1$ if $i > d-1-m$. Then the phase of Ef has the form

$$x \cdot \xi + t(M\xi \cdot \xi), \quad \xi \in \mathbb{R}^{d-1}.$$

We can assume that $m \leq (d-1)/2$ since otherwise we can replace t by $-t$. Note that m is the minimum between the number of positive and negative principal curvatures of $\mathbb{H}$.

We will prove the following.

Theorem 1.1. *Fix $d \geq 4$. Suppose f is supported in $B^{d-1}(0,2)$ and fix $R \geq 1$ and $\varepsilon > 0$. If $m \leq d/2 - 1$ and $p \geq 2(d+2)/d$, then*

$$(1.1) \qquad \|Ef\|_{L^p(B_R)} \leq C_\varepsilon R^\varepsilon \|f\|_{L^p}.$$

When d is even, this result follows from work of Bourgain–Guth ([4]), see Remark 4.3. By Tao's ε-removal argument ([17]), the theorem holds for $p > 2(d+2)/d$ with no loss of R^ε.

There is an alternative proof of Theorem 1.1 that can be found in the recent paper [12] by Hickman and Iliopoulou, which is discussed further below. Indeed, in the cases where $m < \lceil (d-3)/2 \rceil$, the estimates in [12], which generalize the polynomial partitioning method of [10] and [11], are stronger than Theorem 1.1. The main novelty of our approach will be the use of bilinear estimates of S. Lee and Vargas (see below) and an elementary orthogonality estimate in place of k-linear restriction estimates [2] and the ℓ^2 decoupling theorem of Bourgain and Demeter [3]. This gives a somewhat more elementary proof of Theorem 1.1, which matches the best known restriction estimate for the hyperbolic paraboloid in the case where d is odd and $m \geq (d-1)/2 - 1$, and where d is even and $m = d/2 - 1$. Moreover, it was previously unknown if the bilinear estimates we use implied linear estimates in dimension $d \geq 4$. Some of the geometric observations we take advantage of in our proof may also be useful for future work on restriction estimates for $\mathbb{H}$, in particular in the signature 0 case where d is odd and $m = (d-1)/2$. In this case, the Stein–Tomas theorem is the best known restriction estimate for $\mathbb{H}$.

Related results in the literature

In the case $d = 3$, $m = 1$, Theorem 1.1 was proved independently by Vargas ([19]) and S. Lee ([14]) using a bilinear method. This was later improved by Cho and J. Lee ([6]), who adapted the polynomial partitioning method developed by Guth in [9] to show that (1.1) holds for $p > 3.25$. In [15], Stovall proved certain endpoint cases when $d = 3$ that do not follow from arguments in [14] and [19]. See also the paper [13] by Kim. For other recent progress on restriction estimates for perturbations of the hyperbolic paraboloid in dimension 3, see the recent papers of Buschenhenke–Müller–Vargas [5] and Guo–Oh [8].

When $d \geq 4$, the bilinear-to-linear reduction applied by Vargas and S. Lee breaks down for reasons we discuss further in Section 3.1. Improved restriction estimates also do not follow immediately from the techniques established by Guth in [10] to study elliptic paraboloids in higher dimensions. Indeed, the transverse equidistribution estimates that play a crucial role in Guth's argument can fail for hyperbolic paraboloids in certain cases (see Example 8.8 in [11]). For related reasons, the Bourgain–Guth method developed in [4] also does not easily adapt to hyperbolic paraboloids in the case where $d \geq 5$ is odd, although if d is even then the estimate in Theorem 1.1 follows from their more general estimates for Hörmander-type operators (see Remark 4.3)

As mentioned above, Hickman and Iliopoulou [12] have recently extended the polynomial partitioning method developed by Guth in [10] and by Guth, Hickman, and Iliopoulou in [11] to the hyperbolic case. The key new ingredient is the introduction of certain weakened transverse equidistribution estimates. These estimates, which describe the extent to which Ef can concentrate along the neighborhood of a lower-dimensional variety, get worse as the parameter m increases but are still strong enough to obtain improved restriction bounds when m is not too large.

Certain sharp fractal estimates for Ef have also been recently obtained by the author, Erdoğan, and Harris in [1]. These estimates extend the fractal restriction argument of Du

and Zhang [7] to the hyperbolic case. A weighted version of the bilinear argument in this paper plays a key role in the proof of some of the sharp results in the main theorem in [1].

Overview of the paper

The main goal for this paper is to prove Theorem 1.1 using the bilinear restriction estimates proved by S. Lee and Vargas, stated precisely in Theorem 3.1 in Section 3 below. It was previously unknown if these bilinear estimates could be used to prove linear estimates in the range $p \geq 2(d+2)/d$, due to a number of geometric obstructions that arise when trying to apply the usual bilinear-to-linear method in dimension $d \geq 4$. Our argument will follow a broad-narrow scheme adapted from [4], [7], [10]. This broad-narrow analysis allows us to use the estimates of S. Lee and Vargas except in certain exceptional cases which we analyze in Section 3. The main idea is the following: if τ_1 and τ_2 are two caps in the support of f and we do not have a favorable estimate for $Ef_{\tau_1} Ef_{\tau_2}$, then τ_1 and τ_2 must be arranged in a neighborhood of a hyperbolic cone $\mathcal{C}_m$. If we can find no pairs (τ_1, τ_2) for which bilinear estimates apply then the geometry of $\mathcal{C}_m$ forces the caps to in fact be contained in a neighborhood of an m-dimensional plane; we can then treat this scenario using a 'narrow' flat decoupling argument and induction on the scale, at least when $m \leq d/2 - 1$. In the special case where d is odd and $m = (d-1)/2$ our method breaks down since the induction no longer closes. Note however that we always have $m \leq (d-1)/2$.

We review some basic tools that we will use frequently in Section 2. In Section 3 we discuss some more history and background surrounding bilinear restriction estimates. The key lemma describing how bilinear estimates for Ef can fail is then proved in Section 3.3. Finally, in Section 4 we carry out the broad-narrow argument to complete the proof of Theorem 1.1. Some remarks about the failure of our argument in the case where d is odd and $m = (d-1)/2$ can be found at the end of Section 4.

Notation

We will write $A \lesssim B$ if there is some constant $c > 0$ depending only on the dimension and various Lebesgue exponents such that $A \leq cB$. If $A \lesssim B$ and $B \lesssim A$ we also write $A \sim B$. Our uniform constants may also vary from line-to-line, which is allowed since they will remain independent of R. We will also write $A \lessapprox B$ to signify that for each $\varepsilon > 0$ there is some $C(\varepsilon)$ such that $A \leq C(\varepsilon)R^\varepsilon B$. Finally, we will write $A \ll B$ if $A/B \to 0$ as $R \to \infty$.

Let B_r be a ball of radius r in $\mathbb{R}^d$ and let $B_{r^{-1}}$ denote a ball centered at the origin in $\mathbb{R}^d$ of radius r^{-1}. We let w_{B_r} be a smooth weight adapted to B_r in the following sense: $w_{B_r}(x, t)$ decays rapidly for $(x, t) \notin B_r$, and $\widehat{w_{B_r}}$ is supported in a fixed dilate of $B_{r^{-1}}$. Note that we can construct w_{B_r} by taking a bump function w adapted to the unit ball such that

$$|w(x)| \lesssim \frac{1}{(1 + |x|)^{1000d}}$$

and then applying a suitable affine transformation.

If S is a ball or rectangle in $\mathbb{R}^{d-1}$, we let $f_S = f \cdot \phi_S$, where ϕ_S is a smooth bump function supported in a small dilate of S with $\phi_S(\xi) = 1$ when $\xi \in S$. If $\mathcal{M}$ is a smooth manifold and $\rho > 0$, we will let $N_\rho(\mathcal{M})$ denote the ρ-neighborhood of $\mathcal{M}$.

2. Basic tools

In this section we review some basic tools we will use throughout the proof of Theorem 1.1. Below we will always assume that the support of f is contained in $B^{d-1}(0,2)$.

2.1. Wave packet decomposition and parabolic rescaling

We first recall the standard wave packet decomposition for Ef (see for example [6], [10], [14], or [19]). Fix $\rho \in (0,1)$ and suppose $\{\tau\}$ is a collection of finitely-overlapping balls of radius ρ that cover the support of f. We will refer to these τ as ρ-caps. Using a partition of unity we may decompose $f = \sum_\tau f_\tau$, with f_τ supported in a small dilate of τ. Then $Ef = \sum_\tau Ef_\tau$. We let

$$G(\tau) := \frac{(2\xi_1, \ldots, 2\xi_{d-m-1}, -2\xi_{d-m}, \ldots, -2\xi_{d-1}, -1)}{|(2\xi_1, \ldots, 2\xi_{d-m-1}, -2\xi_{d-m}, \ldots, -2\xi_{d-1}, -1)|}$$

when ξ is the center of τ, so $G(\tau)$ is the unit normal direction to $\mathbb{H}$ above the center of τ. If T_τ is any tube in $\mathbb{R}^d$ of dimensions

$$\rho^{-1} \times \cdots \times \rho^{-1} \times \rho^{-2}$$

with long direction $G(\tau)$, then Ef_τ is essentially constant on T_τ.

We also recall that Ef is invariant under parabolic rescalings in the following way.

Proposition 2.1. *Fix $R > 1$ and let $B_R = B^d(0,R)$. Also fix $\rho \in (0,1)$ with $\rho^{-1} < R$. Then for any ρ-cap τ one can find a function g supported in $B^{d-1}(0,2)$ such that*

$$\|g\|_{L^p} = \rho^{-(d-1)/p} \|f_\tau\|_{L^p}$$

and

$$\|Ef_\tau\|_{L^p(B_R)} \leq \rho^{(d-1)-(d+1)/p} \|Eg\|_{L^p(B_{\rho R})}.$$

To prove the proposition, one can use modulation invariance of Ef_τ to reduce to the case where τ is centered at the origin, and then rescale $(x,t) \to (\rho^{-1}\bar{x}, \rho^{-2}\bar{t})$.

The operator Ef has other scaling symmetries that differ from parabolic rescaling, although we will make no use of these symmetries in our arguments. Note, however, that the proof of Theorem 1.1 in the case $d = 3$ due to S. Lee and Vargas ([14], [19]) does exploit these extra symmetries. The same is also true of the Bourgain–Guth proof of the case $d = 3$ (see Remark 4.3 below), along with the improved estimate when $d = 3$ due to Cho and J. Lee in [6].

2.2. Flat decoupling and induction on scales

Decoupling allows us to separate the contribution from different wave packets Ef_τ. This is useful in the 'narrow case' below when we cannot use bilinear restriction estimates. The strongest possible decoupling result for the hyperbolic paraboloid has been proved by Bourgain and Demeter ([3]), though we will not need to use their theorem. Instead it will suffice to use the following more elementary 'flat decoupling' result, which follows easily from orthogonality considerations.

Proposition 2.2 (Flat decoupling). *Suppose that $\mathcal{T}$ is a collection of finitely-overlapping ρ-caps τ with $\rho^{-1} < R$. Then*

$$\Big\| \sum_{\tau \in \mathcal{T}} Ef_\tau \Big\|_{L^p(B_R)} \leq C(\#\mathcal{T})^{1/2-1/p} \Big(\sum_{\tau \in \mathcal{T}} \|Ef_\tau\|_{L^p(w_{B_R})}^2 \Big)^{1/2},$$

where w_{B_R} is a smooth weight adapted to B_R.

Proof. The case $p = \infty$ is just the Cauchy–Schwarz inequality, and when $p = 2$ the proposition follows from Plancharel's theorem. The remaining cases follow by interpolation. $\blacksquare$

Finally, we recall that if R is small enough then Theorem 1.1 follows directly from Hölder's inequality. We can therefore assume by induction that Theorem 1.1 is true at scale ρR whenever $\rho \ll 1$. For technical reasons related to the decoupling result in Proposition 2.2, we also remark that we can assume by induction that the following weighted estimate holds: for any $\varepsilon > 0$,

$$\|Ef\|_{L^p(w_{B_{\rho R}})} \leq C_\varepsilon (\rho R)^\varepsilon \|f\|_{L^p},$$

where $w_{B_{\rho R}}$ is a smooth weight adapted to $B_{\rho R}$.

3. Bilinear restriction estimates for $\mathbb{H}$

In this section we will review some known bilinear estimates and prove a lemma that characterizes what happens if these bilinear estimates fail. The following estimate was proved by S. Lee in dimension $d \geq 3$ (see [14]), and independently by Vargas in dimension 3 (see [19]).

Theorem 3.1 ([14], [19]). *Suppose f_1 and f_2 are supported in open sets τ_1 and τ_2 of diameter ~ 1. If*

$$(3.1) \qquad \inf_{\substack{\xi, \bar{\xi} \in \tau_1 \\ \eta, \bar{\eta} \in \tau_2}} |M(\xi - \eta) \cdot (\bar{\xi} - \bar{\eta})| \geq c > 0,$$

then

$$(3.2) \qquad \||Ef_1 Ef_2|^{1/2}\|_{L^p(B_R)} \leq C_\varepsilon R^\varepsilon \|f_1\|_{L^2}^{1/2} \|f_2\|_{L^2}^{1/2}$$

whenever $p \geq 2(d+2)/d$. If (3.1) fails, then (3.2) can fail as well for all $p \geq 2(d+2)/d$.

We will need to use a version of Theorem 3.1 adapted to K^{-1}-caps for a parameter K such that

$$1 \ll K \ll R.$$

The following is a consequence of Theorem 3.1.

Theorem 3.2. *Suppose f_1 and f_2 are supported in K^{-1}-caps τ_1 and τ_2, respectively, whose centers are separated by CK^{-1}, with C a sufficiently large but uniform constant. Let A be a constant. If*

$$(3.3) \qquad \inf_{\substack{\xi, \bar{\xi} \in \tau_1 \\ \eta, \bar{\eta} \in \tau_2}} |M(\xi - \eta) \cdot (\bar{\xi} - \bar{\eta})| \geq AK^{-1},$$

then

$$(3.4) \qquad \||Ef_1 Ef_2|^{1/2}\|_{L^p(B_R)} \leq C_A K^{O(1)} \|f_1\|_{L^2}^{1/2} \|f_2\|_{L^2}^{1/2}$$

whenever $p \geq 2(d+2)/d$. If (3.3) fails, then (3.4) can fail as well for all $p \geq 2(d+2)/d$.

We say that two K^{-1}-caps τ_1, τ_2 are *strongly separated* if (3.3) holds.

As it is not immediately obvious from scaling that Theorem 3.1 implies Theorem 3.2, we will prove the implication below in Section 3.2.

3.1. Some background

Bilinear restriction estimates in the full range given in Theorem 3.1 were first proved by Wolff in the case of the cone [20]. Wolff's methods were later adapted by Tao in the case of the elliptic paraboloid [16], and then by Vargas and S. Lee independently in the case of hyperbolic paraboloids. In the case of the cone and the elliptic paraboloid, the transversality condition (3.3) is much simpler.

There is an argument due to Tao, Vargas, and Vega ([18]) that allows one to deduce linear restriction estimates from bilinear restriction estimates for elliptic surfaces, and indeed linear restriction estimates are obtained as corollaries of the main results in [16] and [20]. Let E_e denote the extension operator associated to the elliptic paraboloid. The main idea of the argument from [18] is that any two points will belong to a unique pair of dyadic cubes that are separated by a distance proportional to their scale; one can then use this observation to efficiently decompose $|E_e f|^2$ as a sum of terms to which bilinear estimates apply (after a parabolic rescaling). For hyperbolic paraboloids, this argument requires different ideas since the stronger transversality condition (3.3) is more complicated.

In the special case $d = 3, m = 1$ one can apply a simple change variables and instead consider the extension operator associated to the surface

$$\{\xi \in \mathbb{R}^3 : \xi_3 = \xi_1 \xi_2, \ |\xi| \leq 2\}.$$

Then (3.3) is equivalent to the following two-parameter separation condition:

$$(3.5) \qquad |\xi_1 - \eta_1| \gtrsim 1 \quad \text{and} \quad |\xi_2 - \eta_2| \gtrsim 1 \quad \text{for all } \xi \in \tau_1, \eta \in \tau_2.$$

Vargas and S. Lee were able to use this observation to almost recover the bilinear-to-linear reduction from [18], up to certain endpoint cases which were later proved by Stovall [15]. All of these arguments rely on the fact that (3.5) facilitates a two-parameter decomposition of frequency space analogous to the decomposition used in [18]. When $d \geq 4$ this is no longer the case, and the condition (3.3) is no longer well-adapted to Whitney-type decompositions. In particular note that if $d = 3$ then (3.5) can only fail if all the caps

are arranged in a neighborhood of an axis-parallel line (which becomes a diagonal or anti-diagonal line if we undo the change of variables and write the phase as $\xi_1^2 - \xi_2^2$). However, when $d \geq 4$ the estimate (3.4) can fail if the caps are contained near a subset of (a translation of) the hyperbolic cone

$$\mathcal{C} = \{\xi \in \mathbb{R}^{d-1} : \xi_1^2 + \cdots + \xi_{d-m-1}^2 = \xi_{d-m}^2 + \cdots + \xi_{d-1}^2\}.$$

After we deduce Theorem 3.2, we will analyze what can happen in the exceptional case where (3.3) fails for all pairs of caps in the support of f. We will see that failure of (3.3) for every pair of caps forces f to be supported near an affine space of dimension m. We will then be able to use decoupling and induction to prove Theorem 1.1 in the 'narrow' cases where we cannot use Theorem 3.2.

As mentioned in the introduction, our methods do not work when $d = 3, m = 1$. In this case, Theorem 1.1 is still true and follows from arguments by S. Lee, Stovall, or Vargas ([14], [15], [19]). Of course, when $d = 3$, Theorem 1.1 also follows from the stronger restriction estimate due to Cho and J. Lee [6].

3.2. Proof that Theorem 3.1 implies Theorem 3.2

Let e_j denote the standard basis vectors in $\mathbb{R}^{d-1}$. Let τ_1 and τ_2 be two K^{-1}-caps for which (3.3) holds. After translation we can assume that τ_2 is centered at the origin. We may assume that $\text{dist}(\tau_1, \tau_2) \gtrsim K^{-1/2}$ since otherwise the desired result follows easily by rescaling frequency space by $K^{1/2}$.

Since τ_2 is centered at the origin, the condition (3.3) is invariant under linear transformations of the form $U = U' \oplus U''$, where U' is a rotation in $\xi_1, \ldots, \xi_{d-m-1}$ that fixes $\xi_{d-m}, \ldots, \xi_{d-1}$, and U'' is a rotation in $\xi_{d-m}, \ldots, \xi_{d-1}$ that fixes $\xi_1, \ldots, \xi_{d-m-1}$. We can therefore assume that τ_1 is centered at a point of the form

$$\xi^* = (\xi_1, 0, \ldots, 0, \xi_{d-m}, \xi_{d-m+1}, \ldots, \xi_{d-1})$$

with

$$|\xi_1^2 - \xi_{d-m}^2 - \cdots - \xi_{d-1}^2| \geq cK^{-1}.$$

Let us first assume that

$$\xi_1^2 - \xi_{d-m}^2 - \cdots - \xi_{d-1}^2 \geq cK^{-1}.$$

Since we are also assuming $|\xi^*|^2 \geq cK^{-1}$ it follows that

$$(3.6) \qquad \xi_1^2 \geq cK^{-1}.$$

Now let S be the linear transformation such that

$$S\xi^* = e_1, \quad Se_j = e_j, \quad j = 2, \ldots, d-1.$$

One checks using (3.6) that

$$\|S\| \sim \frac{1}{|\xi_1|} \lesssim K^{1/2}.$$

In particular, the first column of S is

$$(1/\xi_1, 0, \ldots, 0, -\xi_{d-m}/\xi_1, \ldots, -\xi_{d-1}/\xi_1),$$

while the other columns are $e_2, \ldots, e_{d-1}$. Now suppose $\eta, \bar{\eta} \in \tau_2$. Since we are assuming that τ_2 is centered at the origin, we then have

$$(3.7) \qquad |M(S\xi^* - S\eta) \cdot (S\xi^* - S\bar{\eta})| = |Me_1 \cdot e_1 + O(K^{-1/2})| \gtrsim 1.$$

Since changing ξ^* to any other $\xi \in \tau_1$ in (3.7) only introduces an error of $O(K^{-1/2})$, it follows that the caps $S\tau_1$ and $S\tau_2$ satisfy the condition (3.1), and so (3.4) follows from (3.2) after rescaling f_1 and f_2 (which is allowed since we can lose $K^{O(1)}$ in the bilinear estimate).

In the case where
$$-\xi_1^2 + \xi_{d-m}^2 + \cdots + \xi_{d-1}^2 \geq cK^{-1},$$
we apply another transformation $U = U' \oplus U''$ to map ξ^* to

$$U\xi^* = (\xi_1, 0, \ldots, 0, \tilde{\xi}_{d-m}, 0, \ldots, 0).$$

Then, since U does not change the norm of either $(\xi_{d-m}, \xi_{d-m+1}, \ldots, \xi_{d-1})$ or $(\xi_1, 0, \ldots, 0)$, it follows that
$$\tilde{\xi}_{d-m} - \xi_1 \geq cK^{-1},$$
and so we repeat the previous argument with $\tilde{\xi}_{d-m}$ playing the role of ξ_1.

3.3. Failure of bilinear estimates

We now prove that if the bilinear estimates in Theorem 3.2 fail, then the caps τ must be localized near an m-dimensional plane. We first prove some geometric lemmas that will lead us in this direction, with the main result of the section being Lemma 3.5 below. Given $\xi = (\xi_1, \ldots, \xi_{d-1}) \in \mathbb{R}^{d-1}$, we will write $\xi' = (\xi_1, \xi_2, \ldots, \xi_{d-m-1})$ and also $\xi'' = (\xi_{d-m}, \ldots, \xi_{d-1})$. The following lemma can be thought of as an approximate polarization identity.

Lemma 3.3. *Let $\mathcal{C}$ denote the surface*

$$\mathcal{C} = \{\xi \in B^{d-1}(0, 2) : \xi_1^2 + \cdots + \xi_{d-m-1}^2 = \xi_{d-m}^2 + \cdots + \xi_{d-1}^2, \ |\xi| < 2\}$$

and let
$$\mathcal{C}_r = \{\xi \in B^{d-1}(0, 2) : |\xi \cdot M\xi| \leq r\}.$$

Suppose $\xi, \eta \in \mathcal{C}_{cK^{-1}}$. Let T_ξ denote the subspace

$$T_\xi = \{\omega \in \mathbb{R}^{d-1} : \omega \cdot M\xi = 0\}.$$

If $\xi - \eta \in \mathcal{C}_{CK^{-1}}$, then η is in an $O(K^{-1})$ neighborhood of T_ξ.

Proof. Since $\xi, \eta \in \mathcal{C}_{cK^{-1}}$ and $\xi - \eta \in \mathcal{C}_{CK^{-1}}$, we have

$$\sum_{i=1}^{d-m-1} (\xi_i - \eta_i)^2 = \sum_{i=d-m}^{d-1} (\xi_i - \eta_i)^2 + O(K^{-1})$$

$$= \sum_{i=1}^{d-m-1} (\xi_i^2 + \eta_i^2) - 2 \sum_{i=d-m}^{d-1} \xi_i \eta_i + O(K^{-1}).$$

After expanding the square on the left side and rearranging, we obtain

$$\xi' \cdot \eta' = \sum_{i=d-m}^{d-1} \xi_i \eta_i + O(K^{-1}) = \xi'' \cdot \eta'' + O(K^{-1}).$$

As a consequence,

$$\eta \cdot M\xi = O(K^{-1}),$$

which proves the lemma. ∎

Lemma 3.4. *Let V be a subspace of $\mathbb{R}^{d-1}$ and suppose that $V \cap B^{d-1}(0, 2) \subset \mathcal{C}_{cK^{-a}}$, where $a > 0$. Then if K is sufficiently large, we must have $\dim V \leq m$.*

Proof. Let $\{v^1, \ldots, v^k\}$ be an orthonormal basis for V. By hypothesis we know that

$$v^i \cdot M v^i = O(K^{-a})$$

for each i. Also note that $v^i - v^j \in V \cap B^{d-1}(0, 2)$ and therefore $v^i - v^j \in \mathcal{C}_{cK^{-a}}$. Then from Lemma 3.3 we conclude that

$$(3.8) \qquad\qquad v^i \cdot M v^j = O(K^{-a})$$

for each pair i, j. Of course,

$$(3.9) \qquad\qquad v^i \cdot v^j = 0, \quad i \neq j,$$

by hypothesis. Now let $P: \mathbb{R}^{d-1} \to \mathbb{R}^m$ denote the orthogonal projection

$$P\omega = (\omega_{d-m}, \ldots, \omega_{d-1}) \in \mathbb{R}^m.$$

From (3.8) and (3.9) we conclude that

$$(3.10) \qquad P v^i \cdot P v^j = O(K^{-a}) \quad \text{if } i \neq j, \quad P v^i \cdot P v^i = \frac{1}{2} + O(K^{-a}).$$

But if K is large enough, depending only on a and the implicit constants above, then (3.10) implies that the set $\{P v^1, \ldots, P v^k\}$ is linearly independent. One way to see this is to note that (3.10) implies that the Gramian matrix G with entries $G_{ij} = \langle P v^i, P v^j \rangle$ is a small perturbation of $\frac{1}{2} I$ when K is large enough, with I the $m \times m$ identity matrix. The claimed

independence follows at once from the characterization of independence in terms of the Gramian matrix. Alternatively, from (3.10) it follows that there is some $\alpha > 0$ such that

$$|\text{Angle}(Pv^i, Pv^j) - \pi/2| \leq \alpha K^{-a}, \quad i \neq j,$$
$$||Pv^i|^2 - 1/2| \leq \alpha K^{-a},$$

which implies the claimed independence if K is large enough (depending on the value of a and α).

Since the elements of $\{Pv^1, \ldots, Pv^k\}$ are all vectors in $\mathbb{R}^m$, we must have $k \leq m$ and so $\dim V \leq m$. ∎

The following lemma is the main result of this section.

Lemma 3.5. *Let $\{\tau\}$ be a collection of finitely-overlapping K^{-1}-caps in $B^{d-1}(0, 2)$ with $Ef = \sum_\tau Ef_\tau$. If K is sufficiently large, then one of the following must occur.*

(i)   *There exist a uniform $\alpha > 0$ and an m-dimensional affine space V such that every τ is contained in an $O(K^{-\alpha})$ neighborhood of V.*

(ii)   *There are two K^{-1}-caps τ and τ' for which*

$$\inf_{\substack{\xi, \bar{\xi} \in \tau \\ \omega, \bar{\omega} \in \tau'}} |M(\xi - \omega) \cdot (\bar{\xi} - \bar{\omega})| \geq AK^{-1}.$$

Proof. Suppose that (ii) fails and let $\tau_0, \tau_1, \ldots, \tau_k$ be distinct caps in $B^{d-1}(0, 2)$ intersecting the support of f. We can assume we can find such caps with $k \geq 2$ or else (i) is trivially true. After modulating Ef, we can also assume that τ_0 is centered at the origin.

Pick $\eta^i \in \tau_i$ for $i = 1, \ldots, k$. Since (ii) fails for each pair of caps (τ_0, τ_i) we see that $\eta^i \in \mathcal{C}_{cK^{-1}}$ for each i, with the constant c depending only on d, A. Since (ii) also fails for each pair (τ_i, τ_j) when $i \neq j$, we see that $\eta^i - \eta^j \in \mathcal{C}_{cK^{-1}}$ as well. Then by Lemma 3.3 we conclude that

$$(3.11) \qquad\qquad\qquad \eta^i \cdot M\eta^j = O(K^{-1})$$

for each i, j (including $i = j$).

Let $\mathcal{S} = \{\eta^1, \ldots, \eta^k\}$. We now construct the space V via a sequence of spaces

$$V_1 \subset V_2 \subset \cdots \subset V_{k'} = V,$$

with

$$V_j = \text{span}\{\eta^{i_1}, \ldots, \eta^{i_j}\}$$

for some subset of $\mathcal{S}$. We note that V will be a vector space since we have shifted τ_0 to the origin. Fix a small parameter $a > 0$, to be determined below. We pick any $\eta^{i_1} \in \mathcal{S}$ with $|\eta^{i_1}| \geq K^{-a}$ and set $V_1 = \text{span}\{\eta^{i_1}\}$. If no such η^{i_1} exists then all of the caps are contained in an $O(K^{-a})$ neighborhood of the origin, and (i) follows with $V = \{0\}$. Now assume, by induction, that we have constructed

$$V_{j-1} = \text{span}\{\eta^{i_1}, \ldots, \eta^{i_{j-1}}\}.$$

If there are any $\eta^{i_j} \in S \setminus \{\eta^{i_1}, \ldots, \eta^{i_{j-1}}\}$ such that

(3.12) $$|\eta^{i_j}| \geq K^{-a} \quad \text{and} \quad \text{Angle}(\eta^{i_j}, V_{j-1}) \geq K^{-a}$$

we pick one and let

$$V_j = \text{span}\{\eta^{i_1}, \ldots, \eta^{i_j}\}.$$

If no such η^{i_j} exists we stop the procedure and let $V = V_{j-1}$. This process continues until there are no more $\eta \in S$ satisfying (3.12). Say this happens at step $k' + 1$, in which case $V = V_{k'}$ (note that there are at most $d - 1$ steps).

After possibly relabeling, we assume that $V = \text{span}\{\eta^1, \eta^2, \ldots, \eta^{k'}\}$. We now claim that $k' \leq m$. This implies (i) since (by construction) the centers of the remaining caps in $B^{d-1}(0, 2)$ make an angle $O(K^{-a})$ with V or are contained in a ball of radius $O(K^{-a})$ centered at the origin. To prove the claim, first note that if $\omega \in V \cap B^{d-1}(0, 2)$ with $\omega = \sum_{i=1}^{k'} a_i \eta^i$, then we have

(3.13) $$\omega \cdot M\omega = \sum_{i,j} a_i a_j (\eta^i \cdot M\eta^j).$$

We claim that it suffices to show that

(3.14) $$|a_i| \lesssim K^{1/2 - \sigma}$$

for some small $\sigma > 0$ depending only on d. Indeed, if (3.14) holds then by (3.13) and (3.11) we have $\omega \cdot M\omega = O(K^{-2\sigma})$, and therefore

$$V \cap B^{d-1}(0, 2) \subset \mathcal{C}_{cK^{-2\sigma}}$$

since ω was arbitrary. But then Lemma 3.4 implies that $k' = \dim V \leq m$ (provided K is sufficiently large), as desired.

We now prove (3.14). After applying an orthogonal transformation, we may assume that V is spanned by the standard basis vectors $\{e_1, \ldots, e_{k'}\}$ and that the j-th component of η^i is 0 for $j > k'$. Without loss of generality we may view ω and the η^i as elements of $\mathbb{R}^{k'}$, since only their first k' components are nonzero. We let B denote the matrix with columns given by the η^i. We also let $B_i(x)$ denote the matrix obtained by replacing the i-th column of B by x (that is, replacing η^i by x). Then we have the following identity, which is essentially Cramer's rule:

(3.15) $$\omega = \frac{1}{\det(B)} \sum_{i=1}^{k'} \det(B_i(\omega)) \eta^i.$$

Indeed, $\det(B_i(\eta^j)) = 0$ if $i \neq j$, so this identity follows by expanding $\det(B_i(\omega))$ using $\omega = \sum_j a_j \eta^j$ and linearity. From (3.15) it follows that $a_i = \det(B_i(\omega))/\det(B)$, and hence, by Hadamard's inequality,

$$|a_i| \leq \frac{1}{|\det(B)|} |\omega| \prod_{j \neq i} |\eta^j| \lesssim \frac{1}{|\det(B)|}.$$

Finally, recall that by construction,

$$(3.16) \qquad \text{Angle}(\eta^i, V_{i-1}) \geq K^{-a} \quad \text{with } V_{i-1} = \text{span}\{\eta^1, \ldots, \eta^{i-1}\},$$

and $|\eta^i| \geq K^{-a}$. Therefore

$$|\det(B)| = |\eta^1 \wedge \cdots \wedge \eta^{k'}| \gtrsim K^{-a(d-1)} K^{-a(d-2)}$$

(using $k' \leq d - 1$). Indeed, the first term comes from rescaling the η^i to have length one, and the second term is a lower bound for the volume of any parallelepiped spanned by unit vectors satisfying (3.16). It follows that $|a_i| \lesssim K^{2a(d-1)}$, and so (3.14) follows if we choose $a < \frac{1}{4(d-1)}$. This completes the proof.

(We remark that the choice of a is far from optimal. For example, taking into account the lengths of the η^i when applying Hadamard's inequality shows that we can take a to be larger than $\frac{1}{4(d-1)}$. However, the precise value of a is not relevant for the proof, so we have chosen to not track it too carefully). ∎

In the next section we will take $K = R^\delta$ for some $\delta = \delta(\varepsilon)$. We are allowed to assume that $K \geq C_\varepsilon$ by induction, and therefore we will always be able to assume K is large enough that Lemma 3.5 applies.

4. The broad-narrow argument

We now prove Theorem 1.1 using a broad-narrow argument adapted from [4], [10], [7]. Fix $\varepsilon > 0$ for the rest of the argument. Let $\delta > 0$ be another parameter with $\delta < \varepsilon^2$ and set

$$K = R^\delta \quad \text{and} \quad K_1 = K^\alpha,$$

where α is as in part (i) of Lemma 3.5. We assume that δ is small enough such that

$$K \ll K_1.$$

Let $\mathcal{T}$ be a collection of finitely-overlapping K^{-1}-caps τ covering the support of f and use a partition of unity to decompose $f = \sum_\tau f_\tau$ with f_τ supported in (a small dilate of) τ. We also let $\{\theta\}$ be a collection of finitely-overlapping K_1^{-1}-caps covering the support of f. Then $f = \sum_\theta f_\theta$ as well.

On the spatial side we fix a collection $\mathcal{Q}$ of finitely-overlapping K^2-cubes that cover $B^d(0, R)$. Given $Q \in \mathcal{Q}$, we define its significant set

$$\mathcal{S}_p(Q) = \left\{ \tau \in \mathcal{T} : \|Ef_\tau\|_{L^p(Q)} \geq \tfrac{1}{100(\#\mathcal{T})} \|Ef\|_{L^p(Q)} \right\}.$$

Note that we have

$$\left\| \sum_{\tau \notin \mathcal{S}_p(Q)} Ef_\tau \right\|_{L^p(Q)} \leq \frac{1}{100} \|Ef\|_{L^p(Q)},$$

and so we will always be able to absorb these error terms into the left-hand side of our estimates for $\|Ef\|_{L^p(Q)}$ below.

Now fix a uniform constant $A > 1$ to be determined below. We say that a K^2-cube Q is *narrow* and write $Q \in \mathcal{N}$ if there is an $(m+1)$-dimensional subspace W such that

$$\text{Angle}(G(\tau), W) \leq AK_1^{-1}$$

for all $\tau \in \mathcal{S}_p(Q)$, where $G(\tau)$ is the unit normal to the surface $\mathbb{H}$ above the center of τ. If a cube Q is not narrow then we say it is *broad* and write $Q \in \mathcal{B}$. We of course have

$$\|Ef\|_{L^p(B_R)}^p \leq \sum_{Q \in \mathcal{N}} \|Ef\|_{L^p(Q)}^p + \sum_{Q \in \mathcal{B}} \|Ef\|_{L^p(Q)}^p,$$

and so it suffices to consider separately the cases when the broad and narrow terms dominate.

4.1. The broad case

We first consider the broad case. We will need to use the following lemma, which is a consequence of Theorem 3.2 and the fact that Ef is essentially constant at scale one. We recall that two caps τ_1 and τ_2 are said to be *strongly separated* if (3.3) holds.

Lemma 4.1. *Suppose f is supported in $B^{d-1}(0, 2)$. Let τ_1 and τ_2 be two strongly separated K^{-1}-caps. Then*

$$\sum_{Q \in \mathcal{B}} \|Ef_{\tau_1}\|_{L^p(Q)}^{p/2} \|Ef_{\tau_2}\|_{L^p(Q)}^{p/2} \leq K^{O(1)} \|f\|_{L^2}^p$$

whenever $p \geq 2(d+2)/d$.

The proof of this lemma is contained in the proof of Proposition 3.1 in [7], though for completeness we include most of the argument.

Proof. We define $f_i(\xi) = e^{ix_i \cdot \xi + t_i \cdot (M\xi \cdot \xi)} f_{\tau_i}(\xi)$ for some choice of $(x_i, t_i) \in \mathbb{R}^d$. Let ϕ be a bump function on $\mathbb{R}^d$ with $\widehat{\phi} = 1$ in $B^d(0, 2)$ and $\widehat{\phi}$ supported in $B^d(0, 3)$. Note that $Ef_i = Ef_i * \phi$ for any choice of (x_i, t_i) in the definition of f_i.

We first fix a single Q. Decompose Q as a union of lattice cubes L_Q of side-length $1/10$. Then we may find (x_i, t_i) as above and $L_Q^* \subset Q$ such that

$$\|Ef_i * \phi\|_{L^\infty(Q)} \leq \|Ef_i * \phi\|_{L^\infty(L_Q^*)}$$

for both $i = 1, 2$. Then

$$\|Ef_{\tau_1}\|_{L^p(Q)}^{1/2} \|Ef_{\tau_2}\|_{L^p(Q)}^{1/2} \leq K^{O(1)} \|Ef_1 * \phi\|_{L^\infty(L_Q^*)}^{1/2} \|Ef_2 * \phi\|_{L^\infty(L_Q^*)}^{1/2}.$$

We may pick our bump function ϕ so that ϕ decays rapidly outside $B^d(0, 1)$ with

$$\sup_{w \in B^d(z,1)} \phi(w) \lesssim \phi(z) \quad \text{for any } z \in \mathbb{R}^d.$$

Therefore,

$$
\|Ef_1 * \phi\|_{L^\infty(L_Q^*)}^{p/2} \|Ef_2 * \phi\|_{L^\infty(L_Q^*)}^{p/2}
$$

$$
\lesssim \left(\int_{L_Q^*} \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} |Ef_1(z_1)| \, |Ef_2(z_2)| \phi(z_1 - z)\phi(z_2 - z) \, dz_1 \, dz_2 \, dz \right)^{p/2}
$$

$$
= C \left(\int_{\mathbb{R}^d} \int_{\mathbb{R}^d} \int_{L_Q^*} |Ef_1(z_1 - z)| \, |Ef_2(z_2 - z)| \phi(z_1)\phi(z_2) \, dz \, dz_1 \, dz_2 \right)^{p/2}.
$$

We now sum over Q. By Minkowsi's and Hölder's inequalities we have

$$
\sum_{Q \in \mathcal{B}} \|Ef_1 * \phi\|_{L^\infty(L_Q^*)}^{p/2} \|Ef_2 * \phi\|_{L^\infty(L_Q^*)}^{p/2}
$$

$$
\lesssim \left[\int_{\mathbb{R}^d} \int_{\mathbb{R}^d} \left(\int_{B_R} |Ef_1(z_1-z)|^{p/2}|Ef_2(z_2-z)|^{p/2}\phi(z_1)^{p/2}\phi(z_2)^{p/2} \, dz \right)^{2/p} dz_1 \, dz_2 \right]^{p/2}
$$

$$
\lesssim \sup_{z_1, z_2} \int_{B_R} |Ef_1(z_1 - z)|^{p/2} \, |Ef_2(z_2 - z)|^{p/2} \, dz
$$

$$
\lesssim \sup_{z_1, z_2} \int_{B_R} |E\widetilde{f}_1(z)|^{p/2} \, |E\widetilde{f}_2(z)|^{p/2} \, dz.
$$

where $\widetilde{f}_i$ is a modulation of f_i that depends on z_i. Note that

$$
\|\widetilde{f}_i\|_{L^2} = \|f_{\tau_i}\|_{L^2}.
$$

Since $\widetilde{f}_i$ is still supported in τ_i and the pair (τ_1, τ_2) is strongly separated, we may apply Theorem 3.2 to conclude that

$$
\sum_{Q \in \mathcal{B}} \|Ef_1 * \phi\|_{L^\infty(L_Q^*)}^{p/2} \|Ef_2 * \phi\|_{L^\infty(L_Q^*)}^{p/2} \leq K^{O(1)} \|f\|_{L^2}^p,
$$

which completes the proof.   ∎

Let Q be a broad cube and first suppose that there is no strongly separated pair of caps in $\mathcal{S}_p(Q)$. Then by Lemma 3.5 there exists an m-dimensional affine space V such that $\tau \subset N_{cK_1^{-1}}(V)$ for all $\tau \in \mathcal{S}_p(Q)$. But this forces the directions $G(\tau)$ to be in an $O(K_1^{-1})$ neighborhood of the $(m + 1)$-plane W in $\mathbb{R}^d$ given by scalar multiples of vectors in $G_0(V)$, where

$$
G_0(\omega) = |(\omega, -1)|G(\omega)
$$

(note that the angle between $G(\omega_1)$ and $G(\omega_2)$ is proportional to the distance $|\omega_1 - \omega_2|$ if the centers of the caps are $O(K^{-1})$-separated). Therefore $Q \in \mathcal{N}$, assuming we have chosen A appropriately depending only on the constant from Lemma 3.5. Since we are assuming $Q \in \mathcal{B}$, this cannot happen, and so there must be two strongly separated caps $\tau_1, \tau_2 \in \mathcal{S}_p(Q)$. By the definition of $\mathcal{S}_p(Q)$ we then have

$$
\|Ef\|_{L^p(Q)} \leq K^{O(1)} \|Ef_{\tau_1}\|_{L^p(Q)}^{1/2} \|Ef_{\tau_2}\|_{L^p(Q)}^{1/2}.
$$

The pair (τ_1, τ_2) depends on Q, but we may make this estimate uniform by summing in ℓ^p over all possible strongly separated pairs (note the number of such pairs is $O(K^{2(d-1)})$). We then apply Lemma 4.1 to conclude that

$$\sum_{Q \in \mathcal{B}} \|Ef\|_{L^p(Q)}^p \le K^{O(1)} \sum_{\substack{(\tau_1, \tau_2) \\ \text{strongly sep.}}} \sum_{Q \in \mathcal{B}} \|Ef_{\tau_1}\|_{L^p(Q)}^{p/2} \|Ef_{\tau_2}\|_{L^p(Q)}^{p/2}$$

$$\le CR^{\varepsilon p} \|f\|_{L^p}^p$$

(provided $\delta = \delta(\varepsilon)$ is chosen small enough, e.g., $\delta = \varepsilon^4$).

4.2. The narrow case

We now estimate the contribution of the narrow cubes. Suppose $Q \in \mathcal{N}$ and let W be an $(m+1)$-plane in $\mathbb{R}^d$ such that

$$\text{Angle}(G(\tau), W) \le A K_1^{-1}$$

for each $\tau \in \mathcal{S}_p(Q)$. Then there is an m-dimensional affine space V in $\mathbb{R}^{d-1}$ such that $\tau \subset N_{cK_1^{-1}}(V)$ for each $\tau \in \mathcal{S}_p(Q)$. In particular, we can take

$$V = \{\omega \in \mathbb{R}^{d-1} : G_0(\omega) \in W\}.$$

We choose a minimal collection Θ_V of θ covering $N_{cK_1^{-1}}(V)$. Note that Θ_V contains cK_1^m caps θ. Applying flat decoupling and then Hölder's inequality we obtain

$$\|Ef\|_{L^p(Q)} \le CK_1^{m(1/2-1/p)} \Big(\sum_{\theta \in \Theta_V} \|Ef_\theta\|_{L^p(w_Q)}^2 \Big)^{1/2}$$

$$\le CK_1^{m(1-2/p)} \Big(\sum_{\theta \in \Theta_V} \|Ef_\theta\|_{L^p(w_Q)}^p \Big)^{1/p} \le CK_1^{m(1-2/p)} \Big(\sum_\theta \|Ef_\theta\|_{L^p(w_Q)}^p \Big)^{1/p}.$$

Since

$$\sum_Q w_Q \lesssim w_{B_R},$$

we can sum over Q to conclude that

$$(4.1) \qquad \Big(\sum_{Q \in \mathcal{N}} \|Ef\|_{L^p(Q)}^p \Big)^{1/p} \le CK_1^{m(1-2/p)} \Big(\sum_\theta \|Ef_\theta\|_{L^p(w_{B_R})}^p \Big)^{1/p}.$$

We will now use induction on scales. By Proposition 2.1, for each θ we can find a function g_θ supported in $B^{d-1}(0, 2)$ such that $\|f_\theta\|_{L^p} = K_1^{-(d-1)/p} \|g_\theta\|_{L^p}$ and such that

$$\|Ef_\theta\|_{L^p(w_{B_R})} \le K_1^{-(d-1)+(d+1)/p} \|Eg_\theta\|_{L^p(w_{B_{R/K_1}})}.$$

By induction on scales we then obtain

$$\|Ef_\theta\|_{L^p(w_{B_R})} \le C_\varepsilon R^\varepsilon K_1^{-\varepsilon} K_1^{-(d-1)+(d+1)/p} K_1^{(d-1)/p} \|f_\theta\|_{L^p}.$$

After applying this argument for each θ, we see from (4.1) that

$$\left(\sum_{Q \in \mathcal{N}} \|Ef\|_{L^p(Q)}^p \right)^{1/p} \leq C_\varepsilon R^\varepsilon K_1^{-\varepsilon} K_1^{m(1-2/p)} K_1^{-(d-1)+(d+1)/p} K_1^{(d-1)/p} \|f\|_{L^p}.$$

The induction closes provided

$$(4.2) \qquad m\left(1 - \frac{2}{p}\right) - (d-1) + \frac{2d}{p} \leq 0,$$

since we may assume K is large enough that $C_\varepsilon K_1^{-\varepsilon} \leq 1$. Note that (4.2) is equivalent to

$$p \geq \frac{2(d-m)}{d-m-1}.$$

Some algebra shows that

$$\frac{2(d-m)}{d-m-1} \leq \frac{2(d+2)}{d}$$

if and only if

$$m \leq \frac{d}{2} - 1.$$

We have assumed this is true for m, and so the narrow case of Theorem 1.1 follows.

Remark 4.2. In the narrow case above we have used flat decoupling in dimension m. This has nothing to do with the curvature of $\mathbb{H}$ and is true for any extension operator $E'f$ when f is supported in a thin neighborhood of an m-plane. If one instead uses the stronger ℓ^2 decoupling result proven by Bourgain and Demeter in [3] there is no gain in our argument, since this still leads to a loss of $K_1^{m(1/2-1/p)}$ in the first step. This is related to the fact that the surface $\mathbb{H}$ contains subsets which are affine spaces of dimension m, even though the curvature of $\mathbb{H}$ is nonzero. The ℓ^2 decoupling does not distinguish the difference, since we can imagine that Ef is supported in a small neighborhood of one of these affine spaces; in this case the $K_1^{m(1/2-1/p)}$ loss is sharp.

We further elaborate on the last claim by considering the special case $d = 5, m = 2$. Note in this case $m = (d-1)/2$ and so our argument in the narrow case does not apply. Fix a K^2-cube Q and suppose there is no pair of caps (τ_1, τ_2) which are strongly separated and in $\mathcal{S}_p(Q)$. Then by Lemma 3.5 the support of f must be contained in an $O(K_1^{-1})$-neighborhood of an m-plane V. If we assume there is at least one significant $\tau \in \mathcal{S}_p(Q)$ that contains the origin, then from the proof of Lemma 3.5 we see that $V \cap B^4(0, 2)$ can be taken to be a subset of the surface $\mathcal{C}$ defined in Section 3. Moreover, V can be assumed to be a vector space.

Let $\{v, u\}$ be an orthonormal basis for V. Since

$$v - u \in V \cap B^4(0,2) \subset \mathcal{C},$$

the argument in Lemma 3.3 implies that $Mv \cdot u = 0$ and hence $Mu \cdot v = 0$. We also know by hypothesis that $Mv \cdot v = 0$ and $Mu \cdot u = 0$. Therefore $\{v, u, Mv, Mu\}$ is an orthonormal basis for $\mathbb{R}^4$ with $V^\perp = \text{span}\{Mv, Mu\}$. Now let A be the orthonormal matrix with inverse

$$A^{-1} = \begin{bmatrix} v & u & Mv & Mu \end{bmatrix},$$

so that A maps V to the 2-plane determined by $\eta_3 = 0$ and $\eta_4 = 0$. Applying the change of coordinates determined by A shows that

$$\|Ef\|_{L^p(Q)} = \|\widetilde{E} f_A\|_{L^p(Q_A)},$$

where f_A is the natural transform of f and $\widetilde{E}$ is the extension operator with phase

$$x \cdot \eta + t(M_A \eta \cdot \eta),$$

where

$$M_A = (A^{-1})^T M A^{-1} = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix}.$$

In particular, $\widetilde{E}$ is the extension operator associated to the hyperbolic surface

$$\mathcal{H} = \{\eta \in \mathbb{R}^5 : \eta_5 = \eta_1 \eta_3 + \eta_2 \eta_4\}.$$

Since f is supported in a K_1^{-1}-neighborhood of V it follows that f_A is supported in a K_1^{-1}-neighborhood of the 2-plane where $\eta_3 = 0, \eta_4 = 0$. As a consequence, $\widehat{\widetilde{E} f_A}$ is supported in a K_1^{-1} neighborhood of the 2-plane

$$V_A = \{\eta \in \mathbb{R}^5 : \eta = (\eta_1, \eta_2, 0, 0, 0)\}.$$

Note that $V_A \subset \mathcal{H}$ and therefore we can choose f so that the loss of $K_1^{2(1/2-1/p)}$ in our first decoupling step is sharp for general f. This can be seen for example by taking f so that $\widehat{\widetilde{E} f_A}$ is essentially the indicator function of $V_A \cap B^5(0, 2)$.

One is tempted to now exploit the non-isotropic scaling symmetry

$$(\eta_1, \eta_2, \eta_3, \eta_4, \eta_5) \to (\eta_1, \eta_2, K_1 \eta_3, K_1 \eta_4, K_1 \eta_5)$$

associated to $\mathcal{H}$ and then argue by induction on scales (since such a transformation will map the support of f_A to a cube of side-length $O(1)$ but shrink the size of Q_A). This gives a favorable result for each individual Q, but remember that V can vary depending on Q and may not even be a vector space. We have not found a way to effectively deal with the contribution of different V, mainly because K_1^{-1}-neighborhoods of different V can intersect in complicated ways and naive estimates give a loss in K_1 that is much too large to close the induction. A similar issue arises in higher dimensions when d is odd and $m = (d - 1)/2$.

Remark 4.3. The idea of using a broad-narrow analysis to deduce linear restriction theorems from multilinear restriction theorems dates back to Bourgain and Guth in [4]. They prove restriction estimates for the (elliptic) paraboloid by using k-linear restriction ([2]) in the broad case and an induction procedure in the narrow case. Their argument works in a range of p that is larger than what Tao proved in [16] using bilinear restriction theorems. When $d = 3$, their methods also adapt to the hyperbolic surface $\mathbb{H}$ and prove Theorem 1.1 in this case. If $d \geq 4$ is even, their methods also prove Theorem 1.1,

and indeed in even dimensions the result follows from their more general estimates for Hörmander-type operators with non-degenerate phases. In this case one can avoid any type of induction-on-dimension procedure in the range $p \geq 2(d + 2)/d$ by directly using the k-linear Bennet–Carbery–Tao estimate with $k = d/2 + 1$, along with a flat decoupling and induction-on-scales argument. In the narrow case in odd dimensions, this procedure is not as effective since one needs to use a smaller k.

Recall that the intersection of $\mathbb{H}$ with a hyperplane can have zero Gaussian curvature. This complicates any induction-on-dimension procedure when compared to the elliptic case, where the intersection of a paraboloid with a hyperplane is a paraboloid of lower dimension. The case $d = 3$ for $\mathbb{H}$ is special since you can only lose curvature if the hyperplane is (almost) parallel to the diagonal $\xi_1 = \xi_2$ or the anti-diagonal $\xi_1 = -\xi_2$. In this case, one can instead exploit non-isotropic scaling symmetries of the operator to close the induction. We have not found a way to carry this argument out in higher dimensions, except in the localized setting summarized at the end of the previous remark. Note that when $d = 3$ there are only two 'bad' directions (the diagonal or anti-diagonal), but when $d \geq 4$ there are infinitely many (any direction along the hyperbolic cone $\mathcal{C}$ defined above). This appears to be one of the key differences between the cases $d = 3$ and $d = 5$, for example.

Remark 4.4. In [12], Hickman and Iliopoulou prove restriction estimates for generalized extension operators with phases which are smooth perturbations of $x \cdot \xi + t(\xi \cdot M\xi)$. It is likely that the bilinear method in this paper will extend to smooth perturbations of the hyperbolic paraboloid $\mathbb{H}$. Indeed, in [14] Lee proves a generalized version of Theorem 3.1 that allows for phases ϕ which are smooth perturbations of $\xi \cdot M\xi$. It is likely that a version of Lemma 3.5 holds, with the plane V replaced by an m-dimensional manifold determined by ϕ. Then the rest of the argument would follow as in Section 4, with only minor changes made. We do not pursue the details here.

References

[1] Barron, A., Erdoğan, M. B. and Harris, T.: Fourier decay of fractal measures on hyperboloids. *Trans. Amer. Math. Soc.* **374** (2021), no. 2, 1041–1075.

[2] Bennet, J., Carbery, A. and Tao, T.: On the multilinear restriction and Kakeya conjectures. *Acta Math.* **196** (2006), no. 2, 261–302.

[3] Bourgain, J. and Demeter, C.: Decouplings for curves and hypersurfaces with nonzero Gaussian curvature. *J. Anal. Math.* **133** (2017), 279–311.

[4] Bourgain, J. and Guth, L.: Bounds on oscillatory integral operators based on multilinear estimates. *Geom. Funct. Analysis* **21** (2011), no. 6, 1239–1295.

[5] Buschenhenke, S., Müller, D. and Vargas, A.: Partitions of flat one-variate functions and a Fourier restriction theorem for related perturbations of the hyperbolic paraboloid. *J. Geom. Anal.* **31** (2021), no. 7, 6941–6986.

[6] Cho, C.-H. and Lee, J.: Improved restriction estimate for hyperbolic surface in $\mathbb{R}^3$. *J. Funct. Anal.* **273** (2017), no. 3, 917–945.

[7] Du, X. and Zhang, R.: Sharp L^2 estimates of the Schrödinger maximal function in higher dimensions. *Ann. of Math. (2)* **189** (2019), no. 3, 837–861.

[8] Guo, S. and Oh, C.: A restriction estimate for surfaces with negative Gaussian curvatures. Preprint 2020, arXiv:2005.12431.

[9] Guth, L.: A restriction estimate using polynomial partitioning. *J. Amer. Math. Soc.* **29** (2016), no. 2, 371–413.

[10] Guth, L.: Restriction estimates using polynomial partitioning II. *Acta Math.* **221** (2018), no. 1, 81–142.

[11] Guth, L., Hickman, J. and Iliopoulou, M.: Sharp estimates for oscillatory integral operators via polynomial partitioning. *Acta Math.* **223** (2019), no. 2, 251–376.

[12] Hickman, J. and Iliopoulou, M.: Sharp L^p estimates for oscillatory integral operators of arbitrary signature. *Math. Z.* **301** (2022), no. 1, 1143–1189.

[13] Kim, J.: Some remarks on Fourier restriction estimates. Preprint 2017, arXiv:1702.01231.

[14] Lee, S.: Bilinear restriction estimates for surfaces with curvatures of different signs. *Trans. Amer. Math. Soc.* **358** (2006), no. 8, 3511–3533.

[15] Stovall, B.: Scale invariant Fourier restriction to a hyperbolic surface. *Anal. PDE* **12** (2019), no. 5, 1215–1224.

[16] Tao, T.: A sharp bilinear restriction estimate for paraboloids. *Geom. Funct. Analysis* **13** (2003), no. 6, 1359–1384.

[17] Tao, T.: The Bochner–Riesz conjecture implies the restriction conjecture. *Duke Math. J.* **96** (1999), no. 2, 363–375.

[18] Tao, T., Vargas, A. and Vega, L.: A bilinear approach to the restriction and Kakeya conjectures. *J. Amer. Math. Soc.* **11** (1998), no. 4, 967–1000.

[19] Vargas, A.: Restriction theorems for a surface with negative curvature. *Math. Z.* **249** (2005), no. 1, 97–111.

[20] Wolff, T.: A sharp bilinear cone restriction estimate. *Ann. of Math. (2)* **153** (2001), no. 3, 661–698.

Alex Barron
Department of Mathematics, University of Illinois, 273 Altgeld Hall,
1409 W. Green Street, Urbana, IL 61801, USA;
aabarron@illinois.edu