Symplectic Geometry and Sheaves

Christopher Kuo

August 14, 2025

Abstract

These are the notes for my talk, *Symplectic Geometry and Sheaves*, presented at the UC Irvine Geometry and Topology Seminar in November 2024. We survey how methods from symplectic geometry refine our understanding of sheaf theory and provide new constructions.

1 Motivation: Why symplectic geometry?

Let M be a smooth manifold. In classical mechanics, one often care about the momenta in addition to the position, and the former transform like covectors, so one has to consider the cotangent bundle T^*M. The "law of physics" governing the trajectories of the particles are often given by an energy function $H \in C^\infty(T^*M)$. For example, H can be the sum of conservative energy and kinetic energy.

As the goal is to understand the trajectory, one has to turn the difference dH to a vector field $X_H \in \mathfrak{X}(T^*M)$, and integrate it. One object to provide such a correspondence is a 2-torsor, i.e., a section ω of $\operatorname{Hom}(T^*(T^*M), T(T^*M)) = (T^*(T^*M))^{\otimes 2}$, which is nondegenerate, and the correspondence is given by

$$\begin{aligned} \Omega^1(T^*M) &\xrightarrow{\sim} \mathfrak{X}(T^*M) \\ v &\mapsto \omega(-, v). \end{aligned}$$

To have a physically meaningful correspondence, this ω has to satisfy the conditions of

1. Energy conserving: The derivative of the energy function by the vector field $X_H(H) = 0$ should be zero. But we have

   $$X_H(H) =: dH(X_H) = \omega(X_H, X_H)$$

   so one can ensure this by assuming ω is alternating. That is, the 2-tensor is a 2-form.

2. Law preserving: Assume X_H integrates to a flow φ_T. Then, $\varphi_t^*(\omega) = \omega$ is constant. By differentiating t and Cartan's formula, we see that

   $$0 = \frac{\partial}{\partial t}\varphi_t^*(\omega) = \varphi_t^*\mathcal{L}_{X_H}(\omega) = \varphi_t^*\left(\cancel{-d(dH)} + (d\omega)(X_H, -)\right).$$

   Thus, this requirement can be implied by the assumption that $d\omega = 0$, i.e., ω is in fact a closed 2-form.

Definition 1.1. A symplectic manifold is a pair (X, ω) where $\omega \in \Omega^2(X)$ is a nondegenerate closed 2-form.

But in addition to its physics origin, there are mathematical reasons why symplectic geometry might be interesting even if the question is about smooth topology. The reason is that the cotangent manifold T^*M admits a canonical 2-form ω_{can} given by the derivative of the Liouville 1-form α_{can}, defined by

$$(\alpha_{can})_{(x,\xi)}(v) \coloneqq \xi \cdot (d\pi_{(x,\xi)} v).$$

A first observation is that any diffeomorphism $M \cong N$ lifts to a symplectomorphism $T^*M \cong T^*N$ with respect to the canonical structure. In fact, any smooth isotopy lifts to a symplectic isotopy. Now, consider the following conjecture by Arnold:

Conjecture 1.2 (Nearby Lagrangian conjecture). Let M be a smooth manifold. If $L \subseteq T^*M$ is a closed exact Lagrangian submanifold, i.e., a half dimensional submanifold such that $\alpha_{can}|_L = df$ for some $f \in C^\infty(L)$, then there exists a symplectic (in fact Hamiltonian) isotopy φ_t moving L to the zero section $\varphi_1(L) = 0_M \subseteq T^*M$.

Proposition 1.3. *Assume the Nearby Lagrangian conjecture is true. Then, the cotangent bundle T^*M detects exotic differential structure on M. More precisely, if M and N are smooth manifold such that there exists an exact symplectomorphism $\psi : T^*M \xrightarrow{\sim} T^*N$, then M and n are diffeomorphic to each other.*

Proof. Since 0_M is a closed exact Lagrangian in T^*M, $\psi(0_M)$ is a closed exact Lagrangian in T^*N. The conjecture implies that there exists a symplectic isotopy φ_t on T^*N such that $\varphi_1(\psi(0_M)) = 0_N$. In other words, the composition $\varphi_1 \circ \psi : M \cong N$ is a diffeomorphism from M to N. □

In the main talk, we will consider a similar paradigm: Questions on the base manifold M can be answered by looking a lifting on T^*M. The objects we study in this talk will, however, be more algebraic.

2 Microlocal sheaf theory

For a topological space X, a sheaf $F \in \mathrm{Sh}(X)$, valued in abelian groups (Ab), is a functor $F : \mathrm{Op}_X^{op} \to (Ab)$, meaning that there are assignments

$$\begin{gathered} U \mapsto F(U) \\ (U \subseteq V) \mapsto (F(V) \to F(U)), \end{gathered}$$

such that the global data are glued from local pieces. More precisely, if U admits an open cover $\{U_i\}_{i \in I}$, then the sequence

$$0 \to F(U) \to \prod_i F(U_i) \to \prod_{i,j} F(U_{ij})$$

should be exact. Here, the first arrow is simply the product of the restrictions map and the second is the difference between two ways of restrictions ($U_i \supseteq U_{ij} \subseteq U_j$, $U_{ij} := U_i \cap U_j$). Indeed, exactness at $F(U)$ means two sections s_1, s_2 in $F(U)$ are the same if their restrictions to U_i are the same for all $i \in I$. Similarly, exactness at $\prod F(U_i)$ means that a family of sections s_i on U_i glues to a section s on U if they agree on double overlaps.

Remark 2.1. For technical reason, we will in fact consider sheaves valued in chain complexes $\mathrm{Ch}(\mathbb{Z})$ (or a suitable stable coefficient), i.e., the target category should be an $(\infty, 1)$-category and one main difference is that the gluing

$$F(U) \to \lim \left(\prod_i F(U_i) \rightrightarrows \prod_{i,j} F(U_{ij}) \Rrightarrow \prod_{i,j,k} F(U_{ijk}) \mathrel{\substack{\to\\[-1ex]\to\\[-1ex]\to\\[-1ex]\to}} \cdots \right)$$

in general does not terminate at any finite step.

Example 2.2. Let M be a manifold. A local system L on M is a representation of $\pi_1(M, x)$ for some fixed point x. To get rid of the choice of x, one has to replace the group by the fundamental groupoid $\pi_1(M)$. Roughly speaking, this is the category consists of points in M and morphisms between two points x and y are given by paths between them. Thus, a local system L is an assignment

$$\begin{aligned} (x \in M) &\mapsto L_x \\ (\gamma : x \sim y) &\mapsto L_x \cong L_y. \end{aligned}$$

In fact, such data combine to a bundle $L \to M$ with discrete fiber and sections of L form a sheaf $\mathcal{L}$ which is locally constant, i.e., $\mathcal{L}(U) \xrightarrow{\sim} \mathcal{L}(V)$ if $V \subseteq U$ are both contractible. Furthermore, all locally constant sheaves are obtained this way so we will identify both notions.

Example 2.3. If $\mathcal{S}$ is a triangulation of M, then any representation $F \in \mathrm{Fun}(\mathcal{S}^{op}, (Ab))$ gives a sheaf. Here, for two simplices X_α and X_β, $\beta \leq \alpha$ if $X_\alpha \subseteq \overline{X_\beta}$. One example is when $M = \mathbb{R}$ and take the triangulation (with open simplices)

$$\mathcal{S} = \{(-\infty, 0), \{0\}, (0, \infty)\}.$$

According to the order, $(-\infty, 0) \leq \{0\}$ so $\mathrm{Fun}(\mathcal{S}^{op}, (Ab))$ is the same as the representations of the quiver

$$\bullet \leftarrow \bullet \rightarrow \bullet.$$

In this case, $\mathbb{Z}_{(0,\infty)}$ the constant sheaf supported by $(0, \infty)$ corresponds to the representation $0 \xleftarrow{\sim} 0 \to \mathbb{Z}$, and similarly $\mathbb{Z}_{[0,\infty)}$ corresponds to $0 \leftarrow \mathbb{Z} \xrightarrow{\sim} \mathbb{Z}$.

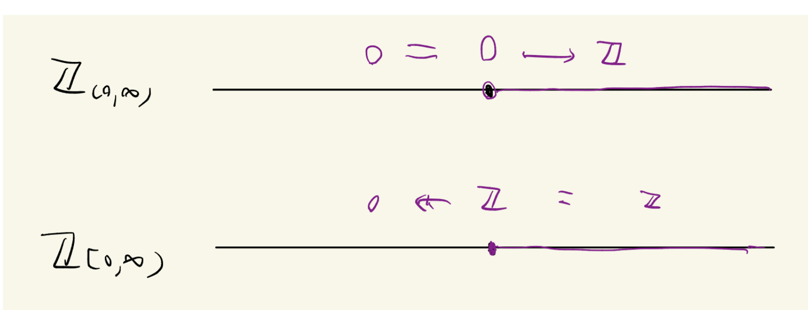

Definition 2.4. For a stratification $\mathcal{S} = \{X_\alpha\}$ of M, a sheaf F is constructible with respect to $\mathcal{S}$ if

$$F|_{X_\alpha} \in \mathrm{Loc}(X_\alpha) \forall \alpha.$$

We use the notation $\mathrm{Sh}_{\mathcal{S}}(M)$ to denote the category of sheaves constructible with respect to $\mathcal{S}$, and $\mathrm{Sh}_{\mathbb{R}-c}(M)$ to denote real constructible sheaves, i.e., the union of $\mathrm{Sh}_{\mathcal{S}}(M)$ such that $\mathcal{S}$ consists of locally closed submanifolds (with certain regularity conditions),

Definition 2.5. The support of F is defined to be $\mathrm{supp}(F) := \overline{\{x | F_x \neq 0\}}$. Equivalently, $\mathrm{supp}(F)^c$ is the largest open set U such that $F|_U = 0$.

Fact 2.6. *If* $F, G \in \mathrm{Sh}(M)$ *such that* $\mathrm{supp}(F) \cap \mathrm{supp}(G) = \varnothing$*, then* $\mathrm{Hom}(G, F) = 0$.

In order to get a more refined measurement, we have to lift it to the cotangent bundle, using the microlocal sheaf theory developed by Kashiwara and Schapira [6].

Remark 2.7. The definitions given so far work both in the abelian and the derived setting. However, beginning from the next definition, we have to assume the coefficient to be $\mathrm{Ch}(\mathbb{Z})$.

Approximate Definition 2.8. For a $F \in \mathrm{Sh}(M)$, its microsupport $\mathrm{SS}(F) \subseteq T^*M$ is a conic closed subset such that $\mathrm{SS}(F) \cap 0_M = \mathrm{supp}(F)$, and for $(x, \xi) \in T^*M \setminus 0_M$, we have $(x, \xi) \notin \mathrm{SS}(F)$ if and only if

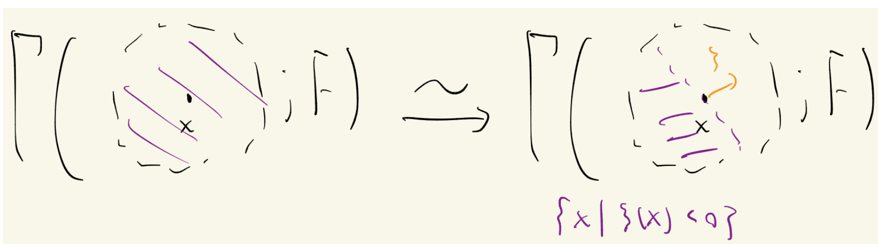

is an equivalence. Roughly speaking, since ξ divides a given open small open ball by half, this says that ξ is not in the microsupport if any section of F propagates uniquely across the barrier dictated by ξ,

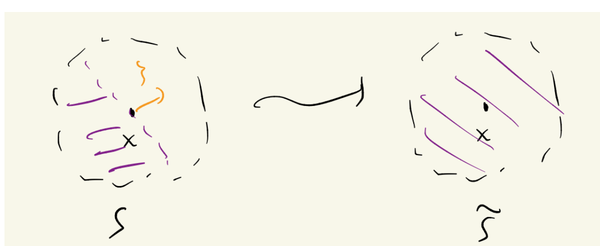

Example 2.9. Both $\mathbb{Z}_{(0,\infty)}$ and $\mathbb{Z}_{[0,\infty)}$ from the earlier examples have the same support, $[0, \infty)$, and are constant away from 0. Thus, it only remains to determine the microsupport at 0. Now, for $\mathbb{Z}_{[0,\infty)}$, there is a section $1 \in \mathbb{Z}$ living on $(0, \infty)$. Since the stalk at 0 is also $\mathbb{Z}$ and the restriction is given by the identity, 1 propagates to the whole $\mathbb{R}$. As a result, $(0, -1) \in$

$T^*\mathbb{R}$ is not in the $SS(\mathbb{Z}_{[0,\infty)})$. The same consideration shows that $(0,1) \in SS(\mathbb{Z}_{[0,\infty)})$ and $SS(\mathbb{Z}_{[0,\infty)}) = T^*_{0,\geq}\mathbb{R}^1 \cup [0,\infty)$. Similarly, one can compute that $SS(\mathbb{Z}_{(0,\infty)}) = T^*_{0,\leq}\mathbb{R}^1 \cup [0,\infty)$.

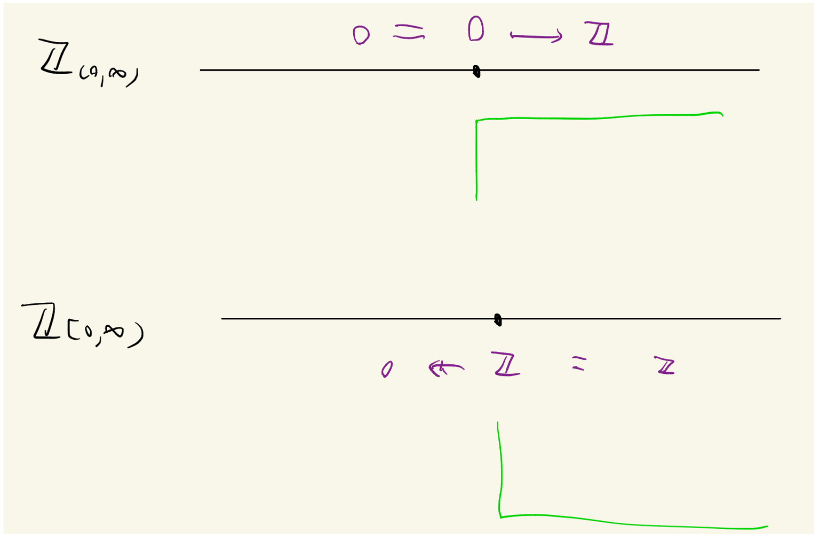

Remark 2.10. In fact, we can show that for any open set $U \subseteq M$ such that ∂U is smooth,

$$SS(\mathbb{Z}_U) = N^*_{out}(U) \text{ and } SS(\mathbb{Z}_{\overline{U}}) = N^*_{in}(U).$$

With this notion, we can upgrade the previous Fact 2.6, which shows that questions regarding the geometry on the base manifold M can be answered by invariants living on the cotangent bundle T^*M. An example of such kind is a generalization that, if L_1 and L_2 are local systems such that L_2 is of perfect stalks, then there is a canonical isomorphism

$$\mathcal{H}\mathrm{om}(L_2, L_1) = L_2^\vee \otimes L_1.$$

Fact 2.11. *If* $SS^\infty(F) \cap SS^\infty(G) = \varnothing$ *and* G *is constructible with perfect stalks, then the canonical morphism*

$$\mathcal{H}\mathrm{om}(G, \mathbb{Z}_M) \otimes F \to \mathcal{H}\mathrm{om}(G, F)$$

is an isomorphism.

Example 2.12. An example where the canonical morphism fails to be isomorphic is the following: Consider $M = \mathbb{R}^2$ with coordinates (x,y), $G = \mathbb{Z}_{\{x,y\geq 0\}}$, and $F = \mathbb{Z}_{\{0\}}$. Then,

$$\mathcal{H}\mathrm{om}(G, \mathbb{Z}_{\mathbb{R}^2}) \otimes F = \mathbb{Z}_{\{x,y>0\}} \otimes \mathbb{Z}_{\{0\}} = 0$$

but

$$\mathcal{H}\mathrm{om}(G, F) = \Gamma_{\{x,y\geq 0\}}\left(\Gamma_{\{0\}}(\mathbb{Z})[2]\right) = \mathbb{Z}_{\{0\}}.$$

3 Symplectic geometry and sheaves

We've seen that thinking microlocally, i.e., thinking locally on the cotangent bundle provides us refined information about sheaves. But to connect the two topics discussed, we recall that T^*M has a canonical symplectic structure given by $d\alpha_{can}$. The source why microlocal sheaf theory is intrinsically symplectic is because of the following theorem by Kashiwara and Schapira:

Theorem 3.1. *Let $F \in \mathrm{Sh}(M)$ be a sheaf.*

1. *The set $\mathrm{SS}(F)$ is coisotropic. If for simplicity we assume $\mathrm{SS}(F)$ is stratified by locally closed submanifolds, this means that the smooth locus $\mathrm{SS}(F)^{sm}$ is a coisotropic submanifold.*

2. *Assume further that M is real analytic and $\mathrm{SS}(F)$ is subanalytic. Then F is constructible if and only if $\mathrm{SS}(F)$ is Lagrangian.*

In fact, starting with Nadler and Zaslow's work [10, 8] and Tamarkin's [12], many results connecting symplectic invariants with sheaf theoretic invariants have been obtained. We'll survey a few later but we begin by mentioning one source of inspiration, a somewhat physical consideration: [1] Take a Lagrangian $L \subseteq T^*M$, with appropriate assumptions such as eventually conic, one can deform it by

$$L_t \coloneqq e^{-t}L, \; t \in \mathbb{R}.$$

Such a Lagrangian defines an algebra via Floer intersection theory. In a similar fashion of Witten deformation [13], this algebra remains constant through out the deformation but, at the limiting situation $t \to +\infty$, the Lagrangian L_∞ becomes conic and its pseudo-holomorphic disks collapse. However, the information is retained in the geometry of the singularity of L_0 remembered by a sheaf.

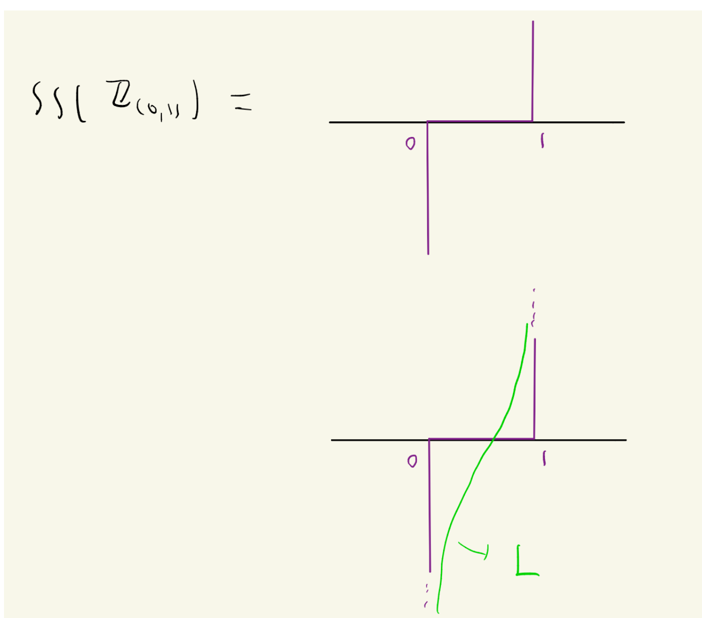

We now mention a few results regarding matching Floer and sheaf-theoretic invariants.

Theorem 3.2 ([3])**.** *Let M be real analytic and $\Lambda \subseteq S^*M$ be a subanalytic Legendrian. Then there is an equivalence*

$$\mathcal{W}(T^*M, \Lambda) = \mathrm{Sh}_\Lambda(M)^c$$

[1] One can find a more detailed discussion in the second arXiv version of [9].

between (the idempotent completion of) the wrapped Fukaya category on T^*M with stop Λ, $\mathcal{W}(T^*M, \Lambda)$, and the compact objects of the category of sheaves microsupported in Λ, $\mathrm{Sh}_\Lambda(M) = \{F | \mathrm{SS}^\infty(F) \subseteq S^*M\}$.

Theorem 3.3 ([4, Theorem E.1])**.** *For a compactly supported Hamiltonian function φ on T^*M, there is an equivalence*

$$\mathrm{HF}^*(\Gamma_\varphi, \Delta_{T^*M}; a, b) = H^*_{M\times M\times[a,b)}\left(\mathrm{Hom}(K_\varphi, \mathbb{Z}_\Delta)\right).$$

Here, HF^* is the Lagrangian Floer homology and a, b means we only the piece of the filtration whose action value is between a and b, and K_φ is certain canonical sheaf induced by φ.

The theorem of Guillermou and Viterbo has the following corollary theorem.

Theorem 3.4 ([7])**.** *Let M be compact without boundary and $U \subseteq T^*M$ be an open subset with a smooth contact boundary ∂U. For $L > 0$, there is an equivalence*

$$\mathrm{SH}^*_{(-\infty,L)}(\overline{U}) = \mathrm{HH}^*(\mathcal{T}(U), T_{L*}).$$

Here SH^* is the symplectic cohomology, which knows about the non-squeezing properties of U, HH^* is the Hochschild homology, and $\mathcal{T}(U)$, the Tamarkin category on U, is a certain enhancement of the category of sheaves.

The above are examples of how microlocal sheaf theory can be applied for the study of symplectic geometry. But we end the notes with the other direction, i.e., how symplectic geometry provides new constructions for sheaf theory. To begin with, we recall the theory of D-modules: Let X be a complex manifold of dimension n, there is a ring-valued sheaf $\mathcal{D}_X$, the ring of differential operators, whose sections are locally of the form

$$P = \sum_{\alpha\in\mathbb{N}^n} c_\alpha \frac{\partial^\alpha}{\partial z^\alpha},\ c_\alpha \in \mathcal{O}_X,\ c_\alpha = 0 \text{ for all but finitely many } \alpha.$$

Here, we choice a local coordinate z_i, and for each multi-index $\alpha \in \mathbb{N}^n$, we denote by

$$\frac{\partial^\alpha}{\partial z^\alpha} = \frac{\partial^{\alpha_1}}{\partial z_1^{\alpha_1}} \cdots \frac{\partial^{\alpha_n}}{\partial z_n^{\alpha_n}}$$

the linear operator given by composition of the standard ones. This ring admits a microlocalization, a ring-valued sheaf $\mathcal{E}_X$ on the coprojective bundle $\mathbb{P}^*X$, whose sections are certain summands of the monomials ξ^α, given by the cotangent coordinates. On the zero section, ξ^α is identified with $\frac{\partial^\alpha}{\partial z^\alpha}$, which allows us to formally invert $\frac{\partial^\alpha}{\partial z^\alpha}$ away from the zero section.

The celebrated Riemann-Hilbert correspondence and its microlocalization, which many, including Kashiwara, Mebkhout, Andronikow, Waschkies, etc., have contributed to, identify a certain class of modules over these rings with constructible sheaf invariants which we've been discussing.

Theorem 3.5. *Let X be a complex manifold. Then there are equivalences*

$$\mathrm{RH} : \mathrm{Perv}(X) = \mathcal{D}_X\text{-}\mathrm{Mod}_{rh}$$

and its microlocalization

$$\mu\mathrm{RH} : \mu\,\mathrm{Perv}(X) = \mathcal{E}_{P^*X}\text{-Mod}_{rh}\,.$$

Here, $\mathrm{Perv}(X)$ *is a certain abelian category inside* $\mathrm{Sh}(X)$*, 'r' standards for regular, and 'h' standards for holonomic.*

Given a complex contact manifold V, i.e., an odd dimensional complex manifold with a maximally non-integrable hyperplane distribution $\mathcal{H} \subseteq TV$. There is a Darboux's theorem in this setting that locally V is contactomorphic some open subsets of $\mathbb{P}^*X$. Kashiwara [5] shows that the sheaf $\mathcal{E}_{P^*X}$ on Darboux charts glues to a canonical global object $\mathcal{E}_V$.

The goal is then to ask whether $\mu\,\mathrm{Perv}$ glues as well and if $\mu\mathrm{RH}$ matches the gluing by both sides. To answer this question, one has to the two questions:

(1) Given a contactomorphism $\chi : \mathcal{U} \cong \mathcal{V}$, between open subsets of some coprojective bundles $\mathbb{P}^*X$ and $\mathbb{P}^*Y$, can $\mu\,\mathrm{Perv}$ be identified?

(2) Given a family of Darboux charts $\{\mathcal{U}_i\}$ how can the identification in (1) be made consistent?

It is already known in [6] that locally if $\mathcal{U}$ and $\mathcal{V}$ are both contractible, then a canonical identification exists. However, compatibility of the gluing data remains unsolved. Indeed, unlike the situation of $\mathcal{E}$, gluing category-valued sheaves a priori requires more data. The question is amazingly resolved by h-principle in contact geometry by Vivek Shende in [11] and later in his joint work with David Nadler [9]. To paraphrase it, the existence question in (1) and (2) are both governed by a homotopical obstruction. With this general machinery, we can show that both $\mu\,\mathrm{Perv}$ and $\mu\mathrm{RH}$ glue.

Theorem 3.6 ([1, 2])**.** *Let V be a complex contact manifold.*

1. *There exists a canonical abelian-category-valued sheaf $\mu\,\mathrm{Perv}_V$ on V, locally isomorphic to $\mu\,\mathrm{Perv}_{\mathbb{P}^*X}$ on Darboux charts.*

2. *There is an equivalence*

$$\mu\mathrm{RH}_V : \mu\,\mathrm{Perv}_V(V) = \mathcal{E}_V\text{-Mod}_{rh}\,.$$

References

[1] Laurent Côté, Christopher Kuo, David Nadler, and Vivek Shende. Perverse microsheaves. *arXiv:2209.12998*, 2022.

[2] Laurent Côté, Christopher Kuo, David Nadler, and Vivek Shende. The microlocal riemann-hilbert correspondence for complex contact manifolds. *arXiv:2406.16222*, 2024.

[3] Sheel Ganatra, John Pardon, and Vivek Shende. Microlocal Morse theory of wrapped Fukaya categories. *arXiv:1809.08807v2*, 2020.

[4] Stéphane Guillermou and Claude Viterbo. The singular support of sheaves is γ-coisotropic. *Geometric and Functional Analysis*, 34(4):1052–1113, 2024.

[5] Masaki Kashiwara. Quantization of contact manifolds. *Publ. Res. Inst. Math. Sci.*, 32(1):1–7, 1996.

[6] Masaki Kashiwara and Pierre Schapira. *Sheaves on manifolds*, volume 292 of *Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1994. With a chapter in French by Christian Houzel, Corrected reprint of the 1990 original.

[7] Christopher Kuo, Vivek Shende, and Bingyu Zhang. On the hochschild cohomology of tamarkin categories. *arXiv:2312.11447*, 2024.

[8] David Nadler. Microlocal branes are constructible sheaves. *arXiv:math/0612399v4*, 2009.

[9] David Nadler and Vivek Shende. Sheaf quantization in Weinstein symplectic manifolds. *arXiv:2007.10154v2*, 2021.

[10] David Nadler and Eric Zaslow. Constructible sheaves and the Fukaya category. *J. Amer. Math. Soc.*, 22(1):233–286, 2009.

[11] Vivek Shende. Microlocal category for Weinstein manifolds via h-principle. *arXiv:1707.07663*, 2017.

[12] Dmitry Tamarkin. Microlocal condition for non-displaceability. *arXiv:0809.1584*, 2008.

[13] Edward Witten. Supersymmetry and Morse theory. *J. Differential Geometry*, 17(4):661–692 (1983), 1982.