

INTRODUCTION TO MICROLOCAL SHEAF THEORY

CHRISTOPHER KUO

ABSTRACT. This document is the notes for V5D4, Introduction to Microlocal Sheaf Theory, taken in Sommersemester 2026 at the University of Bonn. Microlocal sheaf theory is a natural extension of the six-functor formalism in the topological setting to the cotangent bundle: Instead of studying sheaves locally on the base, we refine the analysis of their behavior to be local on the cotangent bundle, using the notion of microsupport, and thus the term “microlocal”.

In the current version, the notes cover the following material: preliminary results from higher category theory, construction of the six-functor formalism using modern machinery, the notion of microsupport and its basic properties, and, lastly, the construction of microsheaves on contact manifolds.

CONTENTS

1. Introduction	2
1.1. From sheaves to microlocal geometry	2
1.2. The course and its goals	2
1.3. Other notable directions	4
Acknowledgement	4
2. Six-functor formalism in topology	5
2.1. Some higher category background	7
2.2. Presentable categories	12
2.3. Anima-valued sheaves	16
2.4. The stable setting	21
2.5. Künneth and duality	26
2.6. The Verdier duality	32
2.7. Category of correspondences	38
2.8. Six-functor formalism in topology: Spectra coefficient	44
2.9. The $*$ -functoriality over general coefficients	50
2.10. Six-functor formalism in topology: General coefficients	54
3. Microlocal sheaf theory	58
3.1. Stalks and support	58
3.2. Microsupport	64
3.3. Compatibility and simple examples	68
3.4. Actions by isotopies	74
3.5. Microlocalization	86
3.6. Globalizing microsheaves	92
References	99

1. INTRODUCTION

1.1. From sheaves to microlocal geometry. Sheaf theory originated as a natural extension of algebraic topology; as Houzel [32] recounts, Jean Leray developed the theory’s foundational concepts while teaching an algebraic topology course at the University of Captivity at Oflag XVII-A in Austria. The question he was then considering concerned the existence of global solutions to nonlinear equations in fluid dynamics; he needed a general framework to bridge the gap between local analysis and global geometry. This procedure of patching local solutions into a global one is now fundamental in modern mathematics and is often referred to as the “local-to-global” principle.

In short, a sheaf on a space X is an assignment F that sends each open set U to a set, group, or ring of sections $F(U)$. Furthermore, for an inclusion $V \subseteq U$, there must be a natural restriction map $F(U) \rightarrow F(V)$, and these maps must respect the topology of X in the same way that functions do: two sections s_1 and s_2 are equal if they agree on an open cover. Moreover, a family of sections $\{s_\alpha\}$, each defined on a member of an open cover $\{U_\alpha\}$, glues to a global section on X provided that, for each pair of indices α and β , the restrictions $s_\alpha|_{U_\alpha \cap U_\beta}$ and $s_\beta|_{U_\alpha \cap U_\beta}$ agree. As one can see, real-valued functions, with possibly varying regularity requirements such as continuity or differentiability, and solutions to a given PDE are all natural examples of sheaves.

Progress, however, did not stop here. Like functions, sheaves can be viewed as geometric objects in their own right, and in many ways they naturally categorify functions. For example, one can pull back functions along maps and, if the functions take values in numbers, one can also multiply them. By categorifying these operations, one can pull back sheaves and, if the sheaves take values in vector spaces, one can also tensor them. Categorification also yields additional structure: there are adjoint notions of pushforward and, in the vector-space-valued case, of internal Hom. In fact, there are two further “exotic” operations that bring the total to six.

The six-functor formalism not only allows us to integrate classical algebraic topology results into a single framework, but it also provides a clean foundation for studying the microlocal aspects of the theory and connecting them with symplectic geometry: over a manifold M , sheaves as geometric objects naturally live not on the base but on the cotangent bundle T^*M . This is realized through the notion of microsupport. For a given sheaf F on M , one can associate a subset $\text{SS}(F) \subseteq T^*M$, the microsupport of F , which records the codirections along which sections do not propagate. This can be further used to construct a sheaf μsh_{T^*M} , valued in categories, of microsheaves; in principle, this assigns to an open set $\Omega \subseteq T^*M$ the category of sheaves, quotiented by those microsupported outside Ω . Even more remarkably, this latter invariant can be globalized using symplectic geometry: for an exact symplectic manifold X , up to the vanishing of a certain obstruction, there exists a sheaf μsh_X locally glued out of μsh_{T^*M} , which is known to contain the wrapped Fukaya category of X .

1.2. The course and its goals. The goal of this course is to give an overview of the current state of the field of microlocal sheaf theory. Our discussion will be divided into two different but intertwining concerns. The first regards the foundations. Kashiwara and Schapira’s tome *Sheaves on Manifolds* [36] encompassed most of the knowledge regarding the subject at the time of its publication and has since served as the standard textbook.

More than thirty years have passed since then. While the book still maintains its prominence both as a textbook and as a source of research directions, many advances in related fields have taken place. First, the higher categorical foundations, developed by many and later consolidated in Lurie’s three books [44, 45, 46] together with his website [47], have transformed the landscape of category theory and homological algebra, which form the foundation upon which microlocal sheaf theory rests. A first goal is therefore to study Chapters 1–3 of [36] in this framework, i.e., to construct the six-functor formalism in topology within this modern setting. We will follow [66], where the material is covered and generalized; for example, many of the restrictions on coefficients and the boundedness assumptions are lifted. Furthermore, a notion of abstract six-functor formalism has been recently established [48, 30, 57], subsuming the one in [36] as a special case.

Our second goal is to discuss the microlocal aspects of the theory, roughly Chapters 4–8 of [36]. The central objects of study here are the microsupport $\mathrm{SS}(F)$, which has acquired a new interpretation through the Ω -lens [29] in the past decade, and the notion of microsheaves $\mu\mathrm{sh}$. Building on these two notions, many of the symplectic aspects of the theory were discovered after the publication of [36]. For example, Guillermou, Kashiwara, and Schapira construct an action of contact isotopies not only on microsheaves but on the categories of sheaves themselves [28], further justifying the intuition of treating constructible sheaves as Lagrangians. Additionally, we will discuss the construction of Nadler and Shende [59, 50], which produces a sheaf of categories $\mu\mathrm{sh}_{X;\tau}$ for any exact symplectic manifold X equipped with a Maslov datum τ .¹

The second part of the course will concern applications. Since microlocal sheaf theory has applications to very different subjects, the exact content will depend on the audience. Here we mention two main applications of the theory.

The first field of application is symplectic geometry. A central question in symplectic geometry is to construct algebraic invariants for Lagrangians; these are often referred to as their quantizations. The classical approach is through Floer theory, a version of infinite-dimensional Morse cohomology. To obtain a more powerful invariant, given a symplectic manifold X , one often collects an admissible family of Lagrangians and forms a Fukaya category $\mathrm{Fuk}(X)$ whose objects are such Lagrangians and whose morphisms, for $L, K \in \mathrm{Fuk}(X)$, are given by the Floer homology $\mathrm{HF}^*(L, K)$.

An alternative approach to quantizing Lagrangians originates with Nadler and Zaslow in [51, 49]. There, they define the infinitesimal Fukaya category $\mathrm{Fuk}_\epsilon(T^*M)$ for a manifold M and show that $\mathrm{Fuk}_\epsilon(T^*M) \simeq \mathrm{Sh}_{\mathbb{R}\text{-c}}(M)$, the category of constructible sheaves. In other words, the Lagrangians they consider admit sheaf quantizations. Related ideas were also pursued by Viterbo [65] and Guillermou [27], as well as by Ganatra, Pardon, and Shende [25, 24, 26], whose work led to the proof of Kontsevich’s conjecture [38].

Another, more classical application is to geometric representation theory. This begins with Beilinson and Bernstein’s localization theorem [6], which identifies $\mathfrak{g}$ -representations—where $\mathfrak{g}$ is the Lie algebra of a complex reductive group G —with D -modules on the flag variety G/B . In the same paper, they also give a proof of the Kazhdan–Lusztig conjecture. One can then further pass from D -modules to perverse sheaves via the Riemann–Hilbert

¹This introduction was written before the beginning of the course. It turned out that the plan was much more ambitious than the author had originally anticipated. The warning here is that we have not covered any of the applications mentioned below. However, we decided to keep the introduction as is, since it could still serve as a guide for further reading.

correspondence [34, 35]; exploiting the simpler structure of perverse sheaves, an independent proof of the conjecture was given by Brylinski and Kashiwara in the same year [10]. A textbook account can be found in [31].

There is a saying, usually attributed to Okounkov [52], that “symplectic resolutions are the Lie algebras of the 21st century.” Indeed, a central notion in geometric representation theory is the category $\mathcal{O}$, and Webster et al. have defined such a notion for symplectic resolutions in [9, 8], using a deformation-theoretic version of D -modules. Building on the work of Andronikof [1] and Waschies [67], and in joint work with Côté at the University of Bonn, Nadler, and Shende, the author was able to microlocalize the Riemann–Hilbert correspondence [12] and provide a description of the category $\mathcal{O}$ of symplectic resolutions in terms of perverse microsheaves.

1.3. Other notable directions. In this section, we list a few other applications of microlocal sheaf theory for the interested reader.²

- (a) Around the same time as Nadler and Zaslow, Tamarkin in [62]³ proves a result concerning non-displaceability: the question of whether two Lagrangians can be displaced by a Hamiltonian isotopy. A related question in contact geometry is the non-squeezing problem [15] or, more quantitatively, the question of constructing capacities [68]. A comparison of the latter with the Floer-theoretic approach can be found in [42].
- (b) One strategy for proving mirror symmetry is to apply [26] and reduce the comparison of quasi-coherent sheaves with Lagrangians to a comparison with constructible sheaves. In the toric case, this is realized as the coherent-constructible correspondence, initiated by Fang, Liu, Treumann, and Zaslow [20]. A complete proof was first given by Kuwagaki [43], with independent proofs around the same time by Vaintrob [64] and, in the proper case, by Zhou [69]. (See [5] for a modern account.) The geometric ingredients needed to apply [26] are obtained in [23].
- (c) There is a technique called persistent homology in the applied mathematics field of topological data analysis, invented to detect noise and small changes that occur abruptly. Kashiwara and Schapira [37] discovered a sheaf-theoretic interpretation of such structure. This notion was subsequently used to define a metric, the interleaving distance, on the category of sheaves by Asano and Ike [2] and by Guillermou and Viterbo [29]. This metric, when restricted to a reasonable subcategory, is complete and can be used to study C^0 symplectic and contact geometry [3, 4].
- (d) Noncommutative geometry pursues the idea of treating a linear category not necessarily with a symmetric product as a kind of “noncommutative” space, and of developing in this setting the usual structures such as Serre duality, spherical functors, and Calabi–Yau structures; it is a particularly useful framework when studying Fukaya categories. By passing to sheaf-theoretic models, one can attempt to realize these structures there. This strategy has been employed by Shende and Takeda [60] to obtain a smooth Calabi–Yau structure. In a series of papers with Li [41, 39, 40], the author gives an independent construction.

Acknowledgement. We thank Yuxin Chen, Yuichi Ike, and Germán Stefanich for their helpful comments, and for reading and pointing out mistakes and typos in the notes.

²We make no claim of exhaustiveness; the author simply discusses here what he is most familiar with.

³arXiv preprint, 2008.

2. SIX-FUNCTOR FORMALISM IN TOPOLOGY

The goal of the first half of Part 1 is to construct the six-functor formalism in topology. Roughly speaking, we want an assignment $X \mapsto \mathrm{Sh}(X)$, from a topological space to its category of sheaves, to have two kinds of functoriality. Let $f: X \rightarrow Y$ be a continuous map. Since sheaves are supposed to categorify functions, or more precisely, cohomology, there should be a notion of pullback $f^*: \mathrm{Sh}(Y) \rightarrow \mathrm{Sh}(X)$. In the cohomological setting, with X and Y nice enough and f proper, one can also push forward along f by “integration along the fiber.” With sheaves, the situation is even better: even without assuming f to be proper, there always exists a proper pushforward $f_!: \mathrm{Sh}(X) \rightarrow \mathrm{Sh}(Y)$. Finally, just as one can multiply two functions when they take values in a ring, when the coefficient category carries a symmetric monoidal product, there is a tensor structure $\otimes$ on $\mathrm{Sh}(X)$. These constitute the basic three functors; the remaining three can be obtained by passing to right adjoints.

Approximate Definition 2.0.1. Fix some symmetric monoidal category $\mathcal{V}$, which one can for simplicity take it as the category of chain complexes $\mathrm{Mod}_{\mathbb{Z}}$. The six-functor formalism is a functor sending a nice topological space X to the category $\mathrm{Sh}(X)$ of $\mathcal{V}$ -valued sheaves on X , equipped with a symmetric monoidal structure $\otimes$ and two kinds of functoriality: for a map $f: X \rightarrow Y$, there is a $*$ -pullback $f^*: \mathrm{Sh}(Y) \rightarrow \mathrm{Sh}(X)$ and a $!$ -pushforward $f_!: \mathrm{Sh}(X) \rightarrow \mathrm{Sh}(Y)$, where $*$ is usually read as “star” and $!$ is read as “shriek.” Furthermore, there are adjunctions

$$f_! \dashv f^!, \quad f^* \dashv f_*, \quad \text{and} \quad F \otimes (-) \dashv \mathcal{H}om(F, -),$$

where $F \in \mathrm{Sh}(X)$ is a sheaf. In addition, they should satisfy the following compatibilities:

- (1) $p_! = p_*$ when p is proper, and $j^! = j^*$ when j is an open embedding.
- (2) $*$ -pullback is symmetric monoidal.
- (3) $!$ -pushforward satisfies the projection formula.
- (4) $*$ -pullback and $!$ -pushforward satisfy base change.

We explain what the compatibilities listed above mean. First, if $f: X \rightarrow Y$ is a map and $F, G \in \mathrm{Sh}(Y)$ are sheaves, then f^* being symmetric monoidal means that there is an isomorphism, functorial in F and G , such that

$$f^*(F \otimes G) \cong (f^*F) \otimes (f^*G),$$

and f^* sends the unit to the unit. In particular, since $\mathrm{Sh}(*) = \mathcal{V}$, if we set $1_X := a^*(1)$, where $a: X \rightarrow \{*\}$ is the projection to a point, then 1_X is the unit of $\mathrm{Sh}(X)$.

For the projection formula, let $F \in \mathrm{Sh}(X)$ and $G \in \mathrm{Sh}(Y)$; then there is an isomorphism

$$f_!(F \otimes f^*G) \cong f_!(F) \otimes G.$$

Finally, to state base change, consider a Cartesian square:

$$\begin{array}{ccc} X' & \xrightarrow{f'} & Y' \\ \downarrow g' & & \downarrow g \\ X & \xrightarrow{f} & Y \end{array}$$

Then, for $F \in \mathrm{Sh}(X)$, there is an isomorphism

$$g^* f_! F \cong f'_! g'^* F.$$

Example 2.0.2 (The Künneth formula). Assume X is a finite CW complex and consider sheaves valued in abelian groups Ab . A classical fact is that, in the derived setting, $\Gamma(X; \mathbb{Z}) \equiv a_*(\mathbb{Z}_X) = C^*(X; \mathbb{Z})$ is the cochain complex of singular cochains. Similarly, $\Gamma_c(X; \mathbb{Z}) \equiv a_!(\mathbb{Z}_X) = C_c^*(X; \mathbb{Z})$ is the complex of compactly supported singular cochains.

Now let X_1 and X_2 be two compact Hausdorff spaces and consider the Cartesian square:

$$\begin{array}{ccc} X_1 \times X_2 & \xrightarrow{p_2} & X_2 \\ \downarrow p_1 & & \downarrow a_2 \\ X_1 & \xrightarrow{a_1} & \{*\} \end{array}$$

Since X_1 and X_2 are compact, all maps in the diagram are proper, so compatibility (1) of Approximate Definition 2.0.1 gives $f_! = f_*$ for each map f in the diagram. We find that the six-functor formalism recovers the Künneth formula:

$$\begin{aligned} \Gamma(X_1; \mathbb{Z}) \otimes \Gamma(X_2; \mathbb{Z}) &= a_{1*}(\mathbb{Z}_{X_1}) \otimes a_{2*}(\mathbb{Z}_{X_2}) \\ &= a_{1*}(\mathbb{Z}_{X_1} \otimes a_1^* a_{2*}(\mathbb{Z}_{X_2})) \\ &= a_{1*}(\mathbb{Z}_{X_1} \otimes p_{1*} p_2^*(\mathbb{Z}_{X_2})) \\ &= a_{1*} p_{1*}(p_1^*(\mathbb{Z}_{X_1}) \otimes p_2^*(\mathbb{Z}_{X_2})) \\ &= a_{1*} p_{1*}(\mathbb{Z}_{X_1 \times X_2}) \\ &= \Gamma(X_1 \times X_2; \mathbb{Z}_{X_1 \times X_2}). \end{aligned}$$

In particular, when both X_1 and X_2 are finite CW complexes, we have

$$C^*(X_1; \mathbb{Z}) \otimes C^*(X_2; \mathbb{Z}) = C^*(X_1 \times X_2; \mathbb{Z}).$$

Example 2.0.3 (The Poincaré duality). Let $f : X \rightarrow Y$ be a map and set $\omega_f := f^!(\mathbb{Z}_Y)$. We claim that, for any $F \in \text{Sh}(Y)$, there is a canonical map $f^*(F) \otimes \omega_X \rightarrow f^!(F)$. Since $f^!$ is a right adjoint, such a map corresponds to a map $f_!(f^*(F) \otimes \omega_X) \rightarrow F$. This latter map is given by the composition

$$f_!(f^*(F) \otimes \omega_X) = F \otimes (f_! \omega_X) = F \otimes (f_! f^! \mathbb{Z}_Y) \rightarrow F \otimes \mathbb{Z}_Y = F.$$

Here, the first equality is the projection formula and the only map that is not an isomorphism is given by the counit $f_! f^! \mathbb{Z}_Y \rightarrow \mathbb{Z}_Y$.

It is a classical fact that, when X and Y are manifolds and $f : X \rightarrow Y$ is smooth, i.e., a submersion, then the map is an isomorphism and ω_f is an invertible local system. In particular, when $Y = \{*\}$, $\omega_X := a^!(\mathbb{Z})$ is called the dualizing sheaf and is concretely given by $\omega_X = \text{or}_X[\dim X]$, the orientation sheaf shifted by the dimension of X . Thus, when X is orientable, $f^! = f^*[\dim X]$. But then, if we further assume that X is compact, we obtain the Poincaré duality:

$$\begin{aligned} C^*(X; \mathbb{Z})^\vee &:= \text{Hom}(a_*(\mathbb{Z}_X), \mathbb{Z}) = \text{Hom}(a_!(\mathbb{Z}_X), \mathbb{Z}) = \text{Hom}(\mathbb{Z}_X, a^! \mathbb{Z}) \\ &= \text{Hom}(a^* \mathbb{Z}, \mathbb{Z}[\dim X]) = a_*(\mathbb{Z}_X)[\dim X] = C^*(X; \mathbb{Z})[\dim X]. \end{aligned}$$

Remark 2.0.4 (The Verdier duality). The object ω_X is first considered by Verdier. Because of the above Example 2.0.3, the morphism

$$\begin{aligned} D_X : \text{Sh}(X) &\rightarrow \text{Sh}(X)^{\text{op}} \\ F &\mapsto \mathcal{H}\text{om}(F, \omega_X) \end{aligned}$$

is often referred to as the Verdier duality and D_X the Verdier dual of F . Despite the name, it is only an isomorphism when restricting to objects with some finiteness condition. However, we will see a modern interpretation⁴ later that does come from an equivalence.

2.1. Some higher category background. We will summarize the needed materials from higher category theory in order to develop sheaf theory. As our goal is to use the theory to talk about things instead of developing it, we will follow the practice of being agnostic of the exact model we use for our ∞ -categories, and refer, for example, to Lurie's work [44] for the exact definition. Before we start, we first mention some motivation of why the classical approach of derived category is not the most convenient.

Example 2.1.1 ([63]). The derived category of an abelian category might not have limits and colimits. Consider for example the bounded derived category $D^b(\mathbb{Z})$ of abelian groups. We first recall that there is a non-trivial map $e : \mathbb{Z}/2 \rightarrow \mathbb{Z}/2[1]$. Indeed,

$$\mathrm{Hom}_{D^b(\mathbb{Z})}(\mathbb{Z}/2, \mathbb{Z}/2[1]) \cong H^1 \mathrm{RHom}(\mathbb{Z}/2, \mathbb{Z}/2) \cong \mathrm{Ext}^1(\mathbb{Z}/2, \mathbb{Z}/2),$$

and the latter can be computed by resolving $\mathbb{Z}/2$ with the free resolution $0 \rightarrow \mathbb{Z} \xrightarrow{2} \mathbb{Z} \rightarrow \mathbb{Z}/2$. Hom-ing it into $\mathbb{Z}/2$, we see that $\mathrm{RHom}(\mathbb{Z}/2, \mathbb{Z}/2)$ can be computed by the chain complex $\mathbb{Z}/2 \xrightarrow{2=0} \mathbb{Z}/2$ so $\mathrm{Ext}^1(\mathbb{Z}/2, \mathbb{Z}/2) = \mathbb{Z}/2$ has exactly one non-zero element. Now the claim is that $K := \ker(e)$ does not exist. Assume otherwise, then we have the following pullback diagram in $D^b(\mathbb{Z})$:

$$\begin{array}{ccc} K & \xrightarrow{i} & \mathbb{Z}/2 \\ \downarrow & & \downarrow e \\ 0 & \longrightarrow & \mathbb{Z}/2[1]. \end{array}$$

Since $H^n(-) \cong \mathrm{Hom}_{D^b(\mathbb{Z})}(\mathbb{Z}[-n], -)$ is limit-preserving for all $n \in \mathbb{Z}$, the same pullback diagram produces a family of short exact sequences

$$0 \rightarrow H^n(K) \xrightarrow{H^n(i)} H^n(\mathbb{Z}/2) \xrightarrow{H^n(e)} H^n(\mathbb{Z}/2[1]) \cong H^{n+1}(\mathbb{Z}/2).$$

This implies that i is a quasi-isomorphism so e must be zero, which is a contradiction, implying that K does not exist.

Example 2.1.2. Let X be a topological space. The abelian category $\mathrm{Loc}(X)$ satisfies descent but its derived category $D_{loc}^b(X)$ doesn't. For example, consider $X = S^2$ and cover it with two open sets U_+ and U_- such that U_+ and U_- are contractible with $U_+ \cap U_- \sim S^1$ homotopic to a circle. As local systems are equivalent to representation of the fundamental group, i.e., $\mathrm{Loc}(X) \cong \mathrm{Fun}(\pi_0(X, x_0), \mathrm{Ab})$, we conclude all three categories $\mathrm{Loc}(S^2) \cong \mathrm{Loc}(U_+) \cong \mathrm{Loc}(U_-) \cong \mathrm{Ab}$ are equivalent to abelian groups. Furthermore, the restriction $\mathrm{Loc}(U_+) \rightarrow \mathrm{Loc}(U_+ \cap U_-)$ is fully faithful. For example,

$$\mathrm{Hom}_{\mathrm{Loc}(S^1)}(\mathbb{Z}_{S^1}, \mathbb{Z}_{S^1}) \cong C^0(S^1; \mathbb{Z}) = \mathbb{Z}.$$

⁴Theorem 2.6.6.

As a result, the diagram of restrictions

$$\begin{array}{ccc} \mathrm{Loc}(S^2) & \xrightarrow{i} & \mathrm{Loc}(U_+) \\ \downarrow & & \downarrow e \\ \mathrm{Loc}(U_-) & \longrightarrow & \mathrm{Loc}(U_+ \cap U_-). \end{array}$$

is a pullback. On the other hand, the same diagram with $\mathrm{Loc}(-)$ replaced by $D_{loc}^b(-)$ is not a pullback. For example, being a pullback in particular implies that morphisms glue locally. However, since $\mathrm{RHom}_{D_{loc}^b(S^2)}(\mathbb{Z}_{S^2}, \mathbb{Z}_{S^2}) \cong C^*(S^2; \mathbb{Z}) \cong \mathbb{Z} \oplus \mathbb{Z}[-2]$, there is a non-trivial map $e : \mathbb{Z}_{S^2} \rightarrow \mathbb{Z}_{S^2}[2]$. On the other hand, since U_+ and U_- are both contractible, $D_{loc}^b(U_+) \cong D_{loc}^b(U_-) \cong D^b(\mathbb{Z})$, so the restriction $e|_{U_+}$ and $e|_{U_-}$ are both 0. This implies that the same diagram with $D_{loc}^b(-)$ is not a pullback, as otherwise we would conclude that $e = 0$.

The problem with the above two examples is that passing to the homotopy categories forgets too much information: in the derived category, two maps $f, g : X^\bullet \rightarrow Y^\bullet$ are the same if and only if they are homotopic, but the homotopy that witnesses their equivalence is a datum that has been thrown away. To maintain the knowledge needed to re-glue the map in Example 2.1.2, one has to include those homotopies back. However, keeping only them will not produce a good theory when needing to identify homotopies, so one has to maintain all the higher homotopies as well.

One solution is then to replace the building blocks of the world with a homotopical version of sets, and one choice are the (∞) -groupoids, categories whose morphisms are all invertible. These objects can be modeled by topological spaces and are thus called spaces in [44]. There are also attempts to treat them as an alternative foundation rather than set theory and, in these frameworks, they are referred to as homotopy types, where the term “type” is used in the sense of mathematical logic [21] instead of topology. Because this lecture is given at the University of Bonn, we will call them *animated sets* or just *anima*.⁵

Notation 2.1.3. As we will take higher category theory as our foundation, we will from now on drop the prefix ∞ and simply refer to ∞ -categories as categories and refer to the usual categories as classical categories. Similarly, we will refer to ∞ -groupoids as groupoids and refer to the usual groupoids as 1-groupoids.

Now, a category $\mathcal{C}$ consists of a collection of objects $\mathrm{Obj}(\mathcal{C})$ and, for a pair of objects X and Y , there is an anima $\mathrm{Hom}_{\mathcal{C}}(X, Y)$ of morphisms from X to Y . This means that arrows of the form $f : X \rightarrow Y$ give “points” in $\mathrm{Hom}_{\mathcal{C}}(X, Y)$. Given two morphisms, $f, g : X \rightarrow Y$, there is an anima $\mathrm{Hom}_{\mathrm{Hom}_{\mathcal{C}}(X, Y)}(f, g)$ of invertible 2-morphisms from f to g . Naively speaking, a “point” $T : f \Rightarrow g$ is an identification, so an invertible arrow, witnessing the “sameness” of f and g , and the process goes on for n -morphisms for any $n \geq 1$. As before, there is a composition map

$$\begin{aligned} \mathrm{Hom}_{\mathcal{C}}(X, Y) \times \mathrm{Hom}_{\mathcal{C}}(Y, Z) &\rightarrow \mathrm{Hom}_{\mathcal{C}}(X, Z) \\ (f, g) &\mapsto g \circ f. \end{aligned}$$

⁵Animation is a general construction that adjoins a certain kind of colimits, the sifted colimits, in the higher categorical sense, to a classical category. For example, the animation of the category of sets gives the ∞ -category of anima, and the animation of k -algebras, for a field k of characteristic 0, gives the ∞ -category of dg k -algebras. See [14, 5.1.4] for a detailed discussion.

But because of the existence of higher morphisms, associativity of composition for example is now structure instead of a property. That is, for three morphisms $f : X \rightarrow Y$, $g : Y \rightarrow Z$, and $h : Z \rightarrow W$, it is part of the data that there is $H \in \text{Hom}_{\text{Hom}_{\mathcal{C}}(X,W)}(h \circ (g \circ f), (h \circ g) \circ f)$ such that

$$H : h \circ (g \circ f) = (h \circ g) \circ f.$$

To relate these different identifications, they are higher identifications, etc. This is often referred to as *coherence structure*.

Example 2.1.4. Consider a topological space X . One can view its points in X as objects in a category. For two points p and q , morphisms $p \rightarrow q$ are given by paths γ from p to q . A 2-morphism between two paths is a homotopy of paths, and so on. Note that, in this case all 1-morphisms are already invertible, and the associated category, by abusing the notation, also denoted by X , is a groupoid.

Even more special, one can consider the case when $X = S^1$ is given by the circle. In this case, since S^1 is path-connected, any two points are the same, but there is a $\mathbb{Z}$ -worth of ways to identify the two points, as the fundamental group $\pi_1(S^1, \{p\}) = \mathbb{Z}$ at any given point $p \in S^1$ is given by the integers. However, the classical result $\pi_n(S^1, \{p\}) = 0$ for $n > 1$ shows that any two such choices are either different or essentially the same. In this language, this is saying that S^1 is in fact the 1-groupoid $B\mathbb{Z}$.

Similar to the situation within categories, the notion of functors also comes with coherent structure now. For example, while a functor still comes with an assignment on objects $\text{Obj}(\mathcal{C}) \rightarrow \text{Obj}(\mathcal{D})$ with a map, for any $X, Y \in \text{Obj}(\mathcal{C})$,

$$F : \text{Hom}_{\mathcal{C}}(X, Y) \rightarrow \text{Hom}_{\mathcal{D}}(F(X), F(Y)),$$

the compatibility with composition, the compatibility $F(g) \circ F(f) = F(g \circ f)$ is now a datum for morphisms $f : X \rightarrow Y$, $g : Y \rightarrow Z$ in $\mathcal{C}$. Similarly so is the associativity, and the identifications between the associativity etc. are now all data that need to be supplied. While this looks cumbersome, keeping track of coherence actually has great advantages. One of them is that there is an abundant supply of well-behaved limits and colimits.

Definition 2.1.5. Let I be a category. We use the notation $I^\triangleright$ to denote the category $I \cup \{\infty\}$ with a final cone point ∞ added to I . A diagram $X_\bullet : I \rightarrow \mathcal{C}$ is said to have a colimit in $\mathcal{C}$ if there exists an extension $\overline{X}_\bullet : I^\triangleright \rightarrow \mathcal{C}$ such that $\overline{X}_\bullet|_I = X_\bullet$ and is initial among all other such extensions. In this case, we write $\overline{X}_\infty := \text{colim}_I X_\bullet$ and call it the limit of the diagram $X_\bullet$. Sometimes, we will also use the notation $\text{colim}_{\alpha \in I} X_\alpha$ or simply $\text{colim}_\alpha X_\alpha$ when it is natural to index by I . A dual notion of limit $\lim_\alpha X_\alpha$ can be defined similarly.

Lemma 2.1.6. *The category Ani admits all small limits and colimits.*

Proof. [47, 02SX]. □

Remark 2.1.7. For a category $\mathcal{C}$, the object $\text{colim}_\alpha X_\alpha$ along with the maps $X_\beta \rightarrow \text{colim}_\alpha X_\alpha$, for $\beta \in I$, satisfies the universal property

$$\text{Hom}_{\mathcal{C}}(\text{colim}_\alpha X_\alpha, Y) = \lim_\alpha \text{Hom}(X_\alpha, Y),$$

for $Y \in \mathcal{C}$, where the limit is now taken in Ani . Thus, it is (homotopically) unique. Similarly, we have

$$\text{Hom}_{\mathcal{C}}(Y, \lim_\alpha X_\alpha) = \lim_\alpha \text{Hom}(Y, X_\alpha)$$

for any $Y \in \mathcal{C}$.

Notation 2.1.8. We denote by Cat the (large) category of small categories with morphisms given by functors.

Remark 2.1.9. The theory we refer to as ∞ -category sits in a larger hierarchy of the theory of (n, m) -category. The number n here denotes the degree at which non-trivial morphisms exist and the number m denotes the degree at which all morphisms beyond it should be invertible.

For example, a classical category has non-trivial morphisms but two morphisms f and g are either different or exactly the same in essentially one way, so it can be viewed as a $(1, 1)$ -category. Another special case is groupoids, where non-trivial morphisms exist for all degrees but they are always invertible, so a groupoid is the same as an $(\infty, 0)$ -category. Their common intersection is a category where two objects X and Y are either strictly the same or different. Thus, up to size issue, a $(1, 0)$ -category is just a set.

What we call an ∞ -category is then an $(\infty, 1)$ -category. The category Cat is most naturally an $(\infty, 2)$ -category, since it is natural to include natural transformations that are not invertible, for example, for the purpose of adjoint functors. We will, however, keep this point implicit.

Proposition 2.1.10. *The category Cat admits all small limits and colimits.*

Proof. [47, Tag 02T0 and Tag 02UN]. □

Remark 2.1.11. Limits in Cat are relatively easy to compute. Let $\mathcal{C}_\alpha : I \rightarrow \text{Cat}$ be a diagram of categories. Then objects in $\mathcal{C} := \lim_{\alpha} \mathcal{C}_\alpha$ are given by compatible families of the form $(x_\alpha)_{\alpha \in I}$ where, for each $\alpha \in I$, x_α is an object of $\mathcal{C}_\alpha$. Given two such families $(x_\alpha)_{\alpha \in I}$, $(y_\alpha)_{\alpha \in I}$, the hom-anima is simply given by the limit hom-anima

$$\text{Hom}_{\mathcal{C}}((x_\alpha)_{\alpha \in I}, (y_\alpha)_{\alpha \in I}) = \lim_{\alpha} \text{Hom}_{\mathcal{C}_\alpha}(x_\alpha, y_\alpha).$$

Unlike limits, colimits are often harder to compute. However, in the case of filtered colimits, there is a simple description.

Proposition 2.1.12. *Let $\mathcal{C}_\alpha : I \rightarrow \text{Cat}$ be a filtered diagram of categories with transition map $f_{\beta\alpha} : \mathcal{C}_\alpha \rightarrow \mathcal{C}_\beta$ and denote by $\mathcal{C} := \text{colim}_{\alpha \in I} \mathcal{C}_\alpha$ their colimit, which comes with a canonical map $f_\alpha : \mathcal{C}_\alpha \rightarrow \mathcal{C}$. Then, objects in $\mathcal{C}$ are given by objects $f_\alpha(x_\alpha)$ for $x_\alpha \in \mathcal{C}_\alpha$. Furthermore, given $x_\alpha \in \mathcal{C}_\alpha$ and $y_\beta \in \mathcal{C}_\beta$, the hom-anima can be computed as*

$$\text{Hom}_{\mathcal{C}}(f_\alpha(x_\alpha), f_\beta(y_\beta)) = \text{colim}_{\alpha, \beta \rightarrow \gamma} \text{Hom}_{\mathcal{C}_\gamma}(f_{\gamma\alpha}(x_\alpha), f_{\gamma\beta}(y_\beta))$$

where γ runs over all indices more final than α and β .

Proof. [55]. □

As a special case of Proposition 2.1.10, the category Cat admits Cartesian product: For $\mathcal{C}, \mathcal{D} \in \text{Cat}$, there exists a product category $\mathcal{C} \times \mathcal{D}$. In fact, it can be organized into a closed symmetric monoidal structure on Cat , usually referred to as the Cartesian monoidal structure [45, 2.4.1]. The internal Hom in this case is given by the functor category: For $\mathcal{C}, \mathcal{D} \in \text{Cat}$, the category $\text{Fun}(\mathcal{C}, \mathcal{D})$ of functors from $\mathcal{C}$ to $\mathcal{D}$ with morphisms given by natural transformations satisfies the property that, for any category $\mathcal{E} \in \text{Cat}$, we have

$$\text{Hom}_{\text{Cat}}(\mathcal{E} \times \mathcal{C}, \mathcal{D}) = \text{Hom}_{\text{Cat}}(\mathcal{E}, \text{Fun}(\mathcal{C}, \mathcal{D})).$$

Similar to the classical situation, the Yoneda lemma is true: Let $\mathcal{C}$ be a category. Consider the assignment $X \mapsto h_X$ where $h_X \in \text{Fun}(\mathcal{C}^{\text{op}}, \text{Ani})$ is the corepresentable functor corepresenting X , $h_X(Y) := \text{Hom}(Y, X)$.

Remark 2.1.13. There is a size issue here since the category Ani is not small; the collection of its objects, containing the collection of all sets, forms a proper class rather than just a set. The usual way around this issue is to fix two universes $\mathcal{U} \subseteq \mathcal{V}$ at the beginning, and declare that a set is small if it lies in $\mathcal{U}$, is large if it lies in $\mathcal{V}$, and it is very large if it lies outside $\mathcal{V}$. This way, Cat consists of those $\mathcal{C}$ such that $\text{Obj}(\mathcal{C}) \in \mathcal{U}$, and for Ani , we mean the category of anima in $\mathcal{U}$, so the category lies in $\mathcal{V}$, and is hence a large category. For this reason, we will consider the very large category $\widehat{\text{Cat}}$ whose objects are large categories.

Definition 2.1.14. To simplify the notation, we will write $\text{PSh}(\mathcal{C}) := \text{Fun}(\mathcal{C}^{\text{op}}, \text{Ani})$ and call it the category of presheaves on $\mathcal{C}$. As discussed in Remark 2.1.13, $\text{PSh}(\mathcal{C}) \in \widehat{\text{Cat}}$ is a large category.

Theorem 2.1.15. *The assignment $(X \mapsto h_X)$ organizes into an embedding $h. : \mathcal{C} \rightarrow \text{PSh}(\mathcal{C})$.*

- (1) *(The strong Yoneda lemma) Assume $\mathcal{C}$ is locally small. That is, $\text{Hom}_{\mathcal{C}}(Y, X)$ is small for all $X, Y \in \mathcal{C}$. Then, for any $F \in \text{PSh}(\mathcal{C})$, evaluation induces the equivalence*

$$\text{Hom}_{\text{PSh}(\mathcal{C})}(h_X, F) = F(X).$$

- (2) *If $\mathcal{C}$ is small, then $h.$ exhibits $\text{PSh}(\mathcal{C})$ as the free cocompletion of $\mathcal{C}$.*

Here (2) means that $\text{PSh}(\mathcal{C})$ admits all small colimits and, for any category $\mathcal{E}$ with all small colimits, the restriction map

$$\text{Fun}^{\text{cocont}}(\text{PSh}(\mathcal{C}), \mathcal{E}) \xrightarrow{\sim} \text{Fun}(\mathcal{C}_0, \mathcal{E})$$

is an equivalence, where we use $\text{Fun}^{\text{cocont}}(-, \mathcal{E}) \subseteq \text{Fun}(-, \mathcal{E})$ to denote the subcategory of functors that preserve all small colimits.

Proof. One can trace the details backwards starting from [47, Tag 03V7]; the strong Yoneda lemma is [47, Tag 03M5]. We only remark that $\text{PSh}(\mathcal{C})$ indeed admits all small limits and colimits. In fact, for any category $\mathcal{D}$ with all small colimits, the category $\text{Fun}(\mathcal{C}^{\text{op}}, \mathcal{D})$ will have all small colimits as they are given pointwise by the formula

$$(\text{colim}_{\alpha} F_{\alpha})(X) := \text{colim}_{\alpha} (F_{\alpha}(X))$$

for any diagram of functors $F. : I \rightarrow \text{Fun}(\mathcal{C}^{\text{op}}, \mathcal{D})$, where the colimit is taken in $\mathcal{D}$, and Lemma 2.1.6 implies that anima is such as category. Existence of limits can be shown similarly, since they are also computed pointwisely. $\square$

Remark 2.1.16. The weak Yoneda lemma is the special case when $F = h_Y$, which implies that

$$\text{Hom}_{\text{PSh}(\mathcal{C})}(h_X, h_Y) = h_Y(X) = \text{Hom}_{\mathcal{C}}(X, Y),$$

so the Yoneda embedding is fully faithful.

Remark 2.1.17. In fact $\text{PSh}(\mathcal{C})$ also admits all small limits, which can be computed pointwise as well.

2.2. Presentable categories. Constructing functors in the higher categorical setting is hard, as it requires specifying infinitely many layers of data. Often, it is easy to begin with a functor that can be constructed classically and use it to build more functors from a standard set of tools. The standard way to begin is through the adjoint functor theorem.

Definition 2.2.1. Let $F : \mathcal{C} \rightarrow \mathcal{D}$ be a functor. A functor $G : \mathcal{D} \rightarrow \mathcal{C}$ is said to be a right adjoint to F if there is a pair of natural transformations $\eta : \text{id}_{\mathcal{C}} \rightarrow G \circ F$ and $\epsilon : F \circ G \rightarrow \text{id}_{\mathcal{D}}$ such that, for any $X \in \mathcal{C}$ and $Y \in \mathcal{D}$, the compositions, depending functorially both on X and Y ,

$$\text{Hom}_{\mathcal{D}}(F(X), Y) \xrightarrow{G} \text{Hom}_{\mathcal{C}}(G(F(X)), G(Y)) \xrightarrow{(-) \circ \eta_X} \text{Hom}_{\mathcal{C}}(X, G(Y))$$

and

$$\text{Hom}_{\mathcal{C}}(X, G(Y)) \xrightarrow{F} \text{Hom}_{\mathcal{D}}(F(X), F(G(Y))) \xrightarrow{\epsilon_Y \circ (-)} \text{Hom}_{\mathcal{D}}(F(X), Y)$$

are inverse to each other. In this case, we write $F \dashv G$ and call η the unit and ϵ the counit. As the definition is symmetric, we also say that F is the left adjoint of G .

Remark 2.2.2. We note that, as being an adjoint is a universal property, once one exists, the functoriality implies that it is unique. Thus, having a left or a right adjoint is a property rather than an extra structure, and, in the above situation, we can write $G = F^R$ or $F = G^L$.

Example 2.2.3. Consider the forgetful map $G : \text{Ab} \rightarrow \text{Set}$, its left adjoint $F : \text{Set} \rightarrow \text{Ab}$ is given by taking the free abelian group $F(X) := \mathbb{Z}^{\oplus X}$ for $X \in \text{Set}$. Denote the basis of $\mathbb{Z}^{\oplus X}$ by e_x , for $x \in X$, the unit is given by

$$\begin{aligned} \eta_X : X &\rightarrow \mathbb{Z}^{\oplus X} \\ x &\mapsto e_x, \end{aligned}$$

and the counit is given by

$$\begin{aligned} \epsilon_M : \mathbb{Z}^{\oplus M} &\rightarrow M \\ e_m &\mapsto m. \end{aligned}$$

Remark 2.2.4. More concretely, the requirement for $F^L \dashv F$ to be an adjunction is that the pair ϵ and η satisfies the triangle identity: For any $X \in \mathcal{C}$ and $Y \in \mathcal{D}$, the composition

$$F(X) \xrightarrow{F(\eta_X)} F F^L F(X) \xrightarrow{\epsilon_{F(X)}} F(X)$$

and

$$F^L(Y) \xrightarrow{\eta_{F^L(Y)}} F^L F F^L(Y) \xrightarrow{F^L(\epsilon_Y)} F^L(Y)$$

are equivalent to identities. In particular, whether a pair of functors F and G is an adjunction can be checked in the homotopy category [54, Remark 4.4.5].

Example 2.2.5 ([47, Tag 02ZZ]). Let $f : \mathcal{C}_0 \rightarrow \mathcal{D}_0$ be a functor between two small categories and $\mathcal{E}$ a category with colimits. Precomposing with f defines a functor

$$(-) \circ f : \text{Fun}(\mathcal{D}_0, \mathcal{E}) \rightarrow \text{Fun}(\mathcal{C}_0, \mathcal{E})$$

between the functor categories. This functor has a left adjoint $\text{Lan}_f(-)$ which on a functor $G : \mathcal{C}_0 \rightarrow \mathcal{E}$ is given by, when evaluating at $Y \in \mathcal{D}_0$,

$$\text{Lan}_f(G)(Y) := \text{colim}_{f(X) \rightarrow Y} G(X)$$

where the index runs over pairs of $X \in \mathcal{C}_0$ and a morphism $f(X) \rightarrow Y$. The $\text{Lan}_f(G) : \mathcal{D}_0 \rightarrow \mathcal{E}$ is called the left Kan extension of G along f . Being a left adjoint, it is initial among $H : \mathcal{D}_0 \rightarrow \mathcal{E}$ that admits a map $G \rightarrow H \circ f$. When $\mathcal{E}$ admits limits, there is a dual notion of right Kan extension $\text{Ran}_f(-)$. Naively on objects, it is given by

$$\text{Ran}_f(G)(Y) := \lim_{Y \rightarrow f(X)} G(X)$$

Often, for simplicity, one would denote the functor $(-) \circ f$ by f^* , the left Kan extension $\text{Lan}_f(-)$ by $f_!$, and $\text{Ran}_f(-)$ by f_* .

Theorem 2.2.6 (The Adjoint Functor Theorem). *Let $\mathcal{C}$ and $\mathcal{D}$ be presentable categories. Then,*

- (1) *A functor $F : \mathcal{C} \rightarrow \mathcal{D}$ is a left adjoint if and only if it preserves colimits.*
- (2) *A functor $G : \mathcal{D} \rightarrow \mathcal{C}$ is a right adjoint if and only if it is accessible and preserves limits.*

Proof. [47, Tag 06Q4]. □

Here, presentability is an assumption on the size of the category. Roughly speaking, we want to mimic Example 2.2.5 and define the right adjoint of F by the formula,

$$F^R(Y) := \lim_{Y \rightarrow F(X)} X$$

so it would be nice for $\mathcal{C}$ to admit limits. But further examining the formula, one realizes that the limit diagram, which takes over all $X \in \mathcal{C}$ that admits a map from Y upon mapping to $\mathcal{D}$ by F , is very likely not small if $\mathcal{C}$ is large enough to have all small limits. The solution is to consider categories large enough to have limits and colimits but is essentially controlled by a small amount of data, so one can reduce the size of the diagram. Such categories are in fact very natural.

Example 2.2.7. We consider an example from classical category theory. Denote by Ab the category of abelian groups. We recall that any abelian group M is a filtered colimit of finitely generated abelian groups, through, for example, the inclusion of finitely generated subgroups.

Furthermore, such finitely generated abelian groups are “small” in the following categorical sense: If M is finitely generated, then for any filtered diagram $\{M_\alpha\}$, the canonical map

$$\text{colim}_\alpha \text{Hom}_{\text{Ab}}(M, M_\alpha) \xrightarrow{\sim} \text{Hom}_{\text{Ab}}(M, \text{colim}_\alpha M_\alpha)$$

is an isomorphism. Indeed, the statement is true for M of the form $\mathbb{Z}^k$, $k \in \mathbb{N}$, but any finitely generated abelian group fits into a short exact sequence $0 \rightarrow \mathbb{Z}^l \rightarrow \mathbb{Z}^k \rightarrow M \rightarrow 0$.

Definition 2.2.8. In this definition, we use $\mathcal{C}$ to denote a category and κ a small regular cardinal.

- (1) ([47, Tag 02P8]) A category I is κ -filtered if any κ -small diagram $K \rightarrow I$ admits an extension to $K^\triangleright \rightarrow I$.
- (2) ([47, Tag 0649]) Assume $\mathcal{C}$ admits κ -filtered colimits. An object $X \in \mathcal{C}$ is κ -compact if the functor $\text{Hom}_{\mathcal{C}}(X, -) : \mathcal{C} \rightarrow \text{Ani}$ preserves κ -filtered colimits. More concretely, if $Y_\alpha : I \rightarrow \mathcal{C}$ is a diagram such that I is κ -filtered, then the canonical map

$$\text{colim}_\alpha \text{Hom}_{\mathcal{C}}(X, Y_\alpha) \xrightarrow{\sim} \text{Hom}_{\mathcal{C}}(X, \text{colim}_\alpha Y_\alpha)$$

is an isomorphism.

- (3) ([47, Tag 0673]) A category $\mathcal{C}$ is κ -compactly generated if every object of $\mathcal{C}$ is a colimit of κ -compact objects.

Remark 2.2.9. Let κ be a cardinal. A set X is said to be κ -small if $|X| < \kappa$. A cardinal κ is regular if for any κ -small index set I and any collection $\{X_\alpha\}_{\alpha \in I}$ of κ -small sets, the union $\bigcup X_\alpha$ is κ -small.

Example 2.2.10. The smallest infinite regular cardinal is aleph zero, $\aleph_0 := |\mathbb{N}|$, as a set X is $\aleph_0$ -small by definition if and only if it is finite. Because of this, one can check that $\aleph_0$ -filtered recovers the usual definition of being filtered. For this reason, we follow the convention of referring to an $\aleph_0$ -compactly generated category simply as a compactly generated category.

Definition 2.2.11 ([47, Tag 06GT]). Let $\mathcal{C}$ be a category.

- (1) We say $\mathcal{C}$ is κ -accessible if it is κ -compactly generated and the full subcategory $\mathcal{C}^\kappa \subseteq \mathcal{C}$ of κ -compact objects is small.
- (2) We say $\mathcal{C}$ is accessible if it is κ -accessible for some small regular cardinal κ .

Remark 2.2.12. We note that there is an inconsistency in the conventions of compactly generated in Example 2.2.10 and accessibility in Definition 2.2.11.

Notation 2.2.13. As in Example 2.2.10, when $\mathcal{C}$ is compactly generated, one often denotes the subcategory of compact objects simply by $\mathcal{C}^c$ or $\mathcal{C}^\omega$ instead of $\mathcal{C}^{\aleph_0}$.

We gave an intrinsic definition for accessibility above. However, it is shown in [47, 06K7] that there is also an extrinsic definition: A category $\mathcal{C}$ is κ -accessible if it is of the form $\text{Ind}_\kappa(\mathcal{C}_0)$, whose definition we now recall for later use.

Definition 2.2.14 ([47, Tag 063J]). For a small category $\mathcal{C}_0$ and small regular cardinal κ , the Ind_κ -completion $\text{Ind}_\kappa(\mathcal{C}_0)$ is obtained by freely joining κ -filtered colimits to $\mathcal{C}_0$. More precisely, there is a fully faithful map $i : \mathcal{C}_0 \hookrightarrow \text{Ind}_\kappa(\mathcal{C}_0)$ such that for any $\mathcal{E}$ with κ -filtered colimits, the restriction map

$$\text{Fun}^{\kappa\text{-}\text{fin}}(\text{Ind}_\kappa(\mathcal{C}_0), \mathcal{E}) \xrightarrow{\sim} \text{Fun}(\mathcal{C}_0, \mathcal{E})$$

is an equivalence. Here we use $\text{Fun}^{\kappa\text{-}\text{fin}}(-, \mathcal{E}) \subseteq \text{Fun}(-, \mathcal{E})$ to denote the subcategory of functors that preserve κ -filtered colimits.

Remark 2.2.15. Let $\mathcal{C}_0$ be a small category. As explained in [47, Tag 04BH], Ind_κ -completion can be constructed as the smallest category in $\text{PSh}(\mathcal{C}_0)$ containing corepresentable functors, the image of $\mathcal{C}_0$ under the Yoneda embedding, and closed under taking κ -filtered colimits. Thus, the filtered colimit $\text{colim}_\alpha h_{X_\alpha}$ taken in $\text{PSh}(\mathcal{C}_0)$, which is often written as the formal filtered colimit “ $\text{colim}_\alpha X_\alpha$ ”, are objects of $\text{Ind}_\kappa(\mathcal{C}_0)$. As explained in [47, Tag 065H], all objects in $\text{Ind}_\kappa(\mathcal{C}_0)$ are of this form. This fact is very useful since we can then compute morphisms more easily. Let “ $\text{colim}_\alpha X_\alpha$ ” and “ $\text{colim}_\beta Y_\beta$ ” be objects of $\text{Ind}_\kappa(\mathcal{C}_0)$. Then

$$\begin{aligned} \text{Hom}_{\text{Ind}(\mathcal{C}_0)}(\text{“colim}_\alpha X_\alpha”, \text{“colim}_\beta Y_\beta”) &= \text{Hom}_{\text{PSh}(\mathcal{C}_0)}(\text{“colim}_\alpha X_\alpha”, \text{“colim}_\beta Y_\beta”) \\ &= \lim_\alpha \text{Hom}_{\text{PSh}(\mathcal{C}_0)}(X_\alpha, \text{“colim}_\beta Y_\beta”) \\ &= \lim_\alpha \text{colim}_\beta \text{Hom}_{\mathcal{C}_0}(X_\alpha, Y_\beta). \end{aligned}$$

Here, the second equality follows from the universal property of the colimit and the last equality is the Yoneda lemma.

Example 2.2.16. Example 2.2.7 implies that the category $\mathbf{Ab}$ is the Ind-completion of the category of finitely generated abelian groups. One can show that similarly, the category of sets $\mathbf{Set}$ is the Ind-completion of finite sets, and the category of rings is the Ind-completion of finitely generated rings.

Example 2.2.17 ([47, Tag 06GV]). The category $\mathbf{Ani}$ is compactly generated and the compact objects are given by finitely dominated anima $\mathbf{Ani}^c$. This is the subcategory generated by the point $*$ $\in \mathbf{Ani}$ under finite colimits and retractions. From the spaces viewpoint, $X \in \mathbf{Ani}^c$ if there is a finite CW-complex K such that X is a retract of K .

The definition of accessibility takes care of the “controlled by a small amount of data” part. We also require the part that $\mathcal{C}$ admits both limits and colimits. In fact, requiring one automatically gives the other.

Definition 2.2.18 ([47, Tag 06NC]). A category $\mathcal{C}$ is presentable if it is accessible and furthermore admits small colimits.

Proposition 2.2.19 ([47, Tag 06PU]). *Assume a category $\mathcal{C}$ is accessible. Then, $\mathcal{C}$ admits small colimits if and only if it admits small limits.*

Corollary 2.2.20. *Let $\mathcal{C}$ be a small category. Then $\mathbf{PSh}(\mathcal{C})$ is presentable.*

Proof. We first observe that $\mathbf{PSh}(\mathcal{C}_0)$, as explained in Theorem 2.1.15, admits small colimits which are computed pointwise, and thus the corepresentable functors are compact objects. Generation by corepresentable functors is a consequence of [47, Tag 03W2]. Naively, for any functor $F \in \mathbf{PSh}(\mathcal{C}_0)$, the canonical map

$$\operatorname{colim}_{h_X \rightarrow T} h_X \xrightarrow{\sim} T$$

can be made precise. (For example, tracing back enough to the proof of [47, Tag 03VX], one can see this by taking the functor $F : \mathcal{C} \rightarrow \mathcal{C}$ to be the Yoneda embedding.) $\square$

Definition 2.2.21 ([47, Tag 06KX]). We say a functor $F : \mathcal{C} \rightarrow \mathcal{D}$ between accessible categories is accessible if F preserves κ -filtered colimits for some κ .

For the purpose of later constructions, we mention that the collection of presentable categories also forms a nice category.

Definition 2.2.22. We denote by $\mathbf{Pr}^L \subseteq \widehat{\mathbf{Cat}}$ the non-full subcategory consisting of presentable categories with morphisms given by colimit-preserving functors. Similarly, we denote by $\mathbf{Pr}^R \subseteq \widehat{\mathbf{Cat}}$ the non-full subcategory with morphisms given by limit-preserving functors which are accessible.

Remark 2.2.23. By Theorem 2.2.6, the $(\infty, \mathbf{1})$ -category $\mathbf{Pr}^L$ (resp. $\mathbf{Pr}^R$) is equivalently given by presentable categories with morphisms being left adjoints (resp. right adjoints). In practice, this latter notion is what we care about. However, it would be a bit funny to define them this way, as being a left (resp. right) adjoint is a statement that only makes sense in an $(\infty, \mathbf{2})$ -category. Still, talking about adjoints is in some sense inevitable; see Section 2.7 for more discussions.

Proposition 2.2.24. *There is an equivalence $(\mathbf{Pr}^L)^{\mathrm{op}} = \mathbf{Pr}^R$ which is given by the identity on objects and by passing to adjoints on morphisms. Furthermore, both $\mathbf{Pr}^L$ and $\mathbf{Pr}^R$ admit all small limits and the inclusions $\mathbf{Pr}^L, \mathbf{Pr}^R \subseteq \widehat{\mathbf{Cat}}$ preserve them.*

Proof. The equivalence $(\Pr^L)^{\text{op}} = \Pr^R$ is explained in [47, Tag 06Q9]. The existence of limits for $\Pr^L$ and $\Pr^R$ and the preservation of them when included in $\widehat{\text{Cat}}$ is explained in [47, Tag 06PH]. $\square$

Remark 2.2.25. Proposition 2.2.24 implies that $\Pr^L$ (and similarly $\Pr^R$) also admits colimits: Consider a diagram $\mathcal{C} : I \rightarrow \Pr^L$ with transition maps given by $F_{\beta\alpha} : \mathcal{C}_\alpha \rightarrow \mathcal{C}_\beta$. The colimit $\text{colim}_I \mathcal{C}$ in $\Pr^L$, using $(\Pr^L)^{\text{op}} = \Pr^R$, can be computed as the limit $\lim_{I^{\text{op}}} \mathcal{C}$ in $\Pr^R$, where we change the diagram to $\mathcal{C} : I^{\text{op}} \rightarrow \Pr^R$ with the same vertices but the morphisms are reversed to $F_{\beta\alpha}^R : \mathcal{C}_\beta \rightarrow \mathcal{C}_\alpha$. In this case, the limit can be computed naively in Cat by Remark 2.1.11.

Example 2.2.26. A funny example is given by the coproduct. Let I be a set, $\mathcal{C}_\alpha$ be a family of presentable categories indexed by I . We know by Proposition 2.2.24 that $\coprod_I \mathcal{C}_\alpha$ in $\Pr^L$ is given by the underlying categories described in Remark 2.1.11, up to cardinality issue. Concretely, this means that $\coprod_I \mathcal{C}_\alpha$ is presentable and the projection map $\prod_{\beta \in I} \mathcal{C}_\beta \rightarrow \mathcal{C}_\alpha$ has a right adjoint.

However, according to Remark 2.2.25, $\coprod_I \mathcal{C}_\alpha$ in $\Pr^L$ is computed by passing to the right adjoints of the functors, since there are none in this case, so it is simply given by $\prod_I \mathcal{C}_\alpha$ in $\Pr^R$ and hence again by the same $\prod_I \mathcal{C}_\alpha$ in Remark 2.1.11. In other words, in $\Pr^L$, the canonical map $\coprod_I \mathcal{C}_\beta \rightarrow \prod_I \mathcal{C}_\alpha$ is an equivalence. Concretely, this means that $\prod_{\beta \in I} \mathcal{C}_\beta \rightarrow \mathcal{C}_\alpha$ also has a left adjoint.

Warning 2.2.27. Lurie changed the meaning of the term “compactly generated” between [44] and [47]. More precisely, in [44, Definition 5.5.7.1], a κ -compactly generated category is a presentable category that is κ -accessible. However, as seen in our discussion, the definition we follow (Definition 2.2.8), taken from [47], is a weaker notion than κ -accessibility.

2.3. Anima-valued sheaves. Our goal here is to define the notion of sheaves valued in anima, show that it is a presentable category, and study the $*$ -adjunction. We mention that in this “unstable” setting⁶, the $!$ -functoriality does not exist. Nevertheless, many important properties, proper base-change for example, already exist in this setting.

Definition 2.3.1. Let X be a topological space. Denote by Op_X the poset of opens in X given by inclusion. For a category $\mathcal{C}$ ⁷, viewed as coefficient, we use the notation $\text{PSh}(X; \mathcal{C}) := \text{Fun}(\text{Op}_X^{\text{op}}, \mathcal{C})$ and call its objects presheaves. When $\mathcal{C} = \text{Ani}$, we would use the notation $\text{PSh}(X) := \text{PSh}(\text{Op}_X) := \text{Fun}(\text{Op}_X^{\text{op}}, \text{Ani})$ for simplicity.

Geometrically, a presheaf F on X is an assignment sending an open set U to an object $F(U) \in \mathcal{C}$ and inclusion $V \subset U$ to a restriction map $F(U) \rightarrow F(V)$. As we are in the higher categorical setting, a double inclusion $W \subseteq V \subseteq U$ is assigned to a two morphism

$$\begin{array}{ccc} F(U) & \xrightarrow{\quad} & F(W) \\ & \searrow & \uparrow \\ & & F(V) \end{array}$$

⁶The term “unstable” is used to emphasize that the situation is different from the stable setup which we will recall in the next Section 2.4.

⁷The “mental drawback” of working in the higher categorical setting is that categories are both the framework and the objects of interest themselves. Thus, I will use the chancery version $\mathcal{C}$ for the first case and the script version $\mathcal{C}$ for the latter.

and so on. To detect the topological information of the space, we would like it to respect the topology.

Definition 2.3.2. Assume $\mathcal{C}$ has a final object. A presheaf $F : \text{Op}_X^{\text{op}} \rightarrow \mathcal{C}$ is a sheaf if it satisfies the following *local-to-global* property: For any open $U \subseteq X$ and any open cover $\{U_\alpha\}_{\alpha \in I}$, one can form an augmented Čech complex,

$$F(U) \rightarrow \prod_{\alpha \in I} F(U_\alpha) \rightrightarrows \prod_{\alpha, \beta \in I} F(U_\alpha \cap U_\beta) \Rrightarrow \prod_{\alpha, \beta, \gamma \in I} F(U_\alpha \cap U_\beta \cap U_\gamma) \Rrightarrow \cdots$$

where the arrows are given by the restrictions, and we require $F(U)$ to be the limit

$$(1) \quad F(U) \xrightarrow{\sim} \lim \left(\prod_{\alpha \in I} F(U_\alpha) \rightrightarrows \prod_{\alpha, \beta \in I} F(U_\alpha \cap U_\beta) \Rrightarrow \prod_{\alpha, \beta, \gamma \in I} F(U_\alpha \cap U_\beta \cap U_\gamma) \Rrightarrow \cdots \right).$$

Often, we would simply write it as $F(U) \xrightarrow{\sim} \lim_{\alpha \in I} F(U_\alpha)$. We denote by $\text{Sh}(X; \mathcal{C}) \subseteq \text{PSh}(X; \mathcal{C})$ the subcategory of sheaves. Similar to the case of presheaves, when $\mathcal{C} = \text{Ani}$, we will use the notation $\text{Sh}(X) := \text{Sh}(X; \text{Ani})$ for simplicity.

Remark 2.3.3. In essence, condition (1) is saying that $F : \text{Op}_X^{\text{op}} \rightarrow \text{Ani}$ is a sheaf if any “compatible” family $\{s_\alpha\}_{\alpha \in I}$ of sections $s_\alpha \in F(U_\alpha)$ should uniquely determine a section. Here, “compatible” should mean that on double overlaps $U_{\alpha\beta}$, we should have $s_\alpha|_{U_{\alpha\beta}} = s_\beta|_{U_{\alpha\beta}}$. However, this is now a structure $h_{\beta\alpha} : s_\alpha|_{U_{\alpha\beta}} = s_\beta|_{U_{\alpha\beta}}$ so on triple overlaps $U_{\alpha\beta\gamma}$, we should have a 2-morphism $c_{\gamma\beta\alpha} : h_{\gamma\beta} \circ h_{\beta\alpha} \xrightarrow{\sim} h_{\gamma\alpha}$ with data of compatibility on quadruple overlaps and so on.

Remark 2.3.4. Note that the having a sheaf F must force $\mathcal{C}$ to have final object $*$. This is because $\emptyset \subseteq X$ is an open set and it can be cover by the empty family. But then, the empty product, if exists in $\mathcal{C}$, has to be a final object $*$ in $\mathcal{C}$, and the existence of F implies that

$$F(\emptyset) \xrightarrow{\sim} (* \rightrightarrows * \Rrightarrow \cdots) = *.$$

Example 2.3.5. Consider the case when $\mathcal{C}$ is the classical category of $\mathbb{R}$ -vector spaces $\text{Vect}_{\mathbb{R}}^{\heartsuit}$. In this case, the complex in Definition 2.3.2 truncates at degree two and linearity further simplifies the sheaf condition to requiring the sequence

$$0 \rightarrow F(U) \xrightarrow{r_\alpha} \prod_{\alpha \in I} F(U_\alpha) \xrightarrow{r_{\alpha\beta, \alpha} - r_{\beta\alpha, \beta}} \prod_{\alpha, \beta \in I} F(U_{\alpha\beta})$$

to be exact. Here we use the notation $U_{\alpha\beta} = U_\alpha \cap U_\beta$.

More precisely, exactness at $F(U)$ means that two sections s_1, s_2 are equal if $s_1|_{U_\alpha} = s_2|_{U_\alpha}$ for all α , and exactness at $\prod_{\alpha} F(U_\alpha)$ means that if a family of sections s_α on U_α agrees on overlaps, i.e., $s_\alpha|_{U_{\alpha\beta}} = s_\beta|_{U_{\alpha\beta}}$, then they come from a global section s on X with $s_\alpha = s|_{U_\alpha}$.

Standard examples of $\text{Vect}_{\mathbb{R}}^{\heartsuit}$ sheaves are given by functions. For example, $\mathbb{R}$ -valued continuous function $C_X^0(U) := C^0(U; \mathbb{R})$ with the obvious restriction is such a sheaf. Similarly, if X , for example, has a smooth/real analytic/holomorphic structure, then the assignment to smooth/real analytic/holomorphic functions $C_X^\infty / C_X^\omega / \mathcal{O}_X$ will be a subsheaf of C_X^0 . Another more homotopical example is the locally constant function $C_X^{lc}(U) := C^{lc}(U; \mathbb{R}) \cong C^0(U; \mathbb{R}^\delta)$, where $\mathbb{R}^\delta$ is the topological space $\mathbb{R}$ with the discrete topology, and one recalls that in good cases it computes the homology.

We will use the inclusion $\mathrm{Sh}(X) \hookrightarrow \mathrm{PSh}(X)$ to deduce that $\mathrm{Sh}(X)$ is also presentable. The procedure holds more generally: localizing along a family of morphisms in a presentable category will produce a presentable category.

Definition 2.3.6 ([47, Tag 06UU]). Let $\mathcal{C}$ be a presentable category. We say that a functor $F : \mathcal{C} \rightarrow \mathcal{D}$ is a Bousfield localization functor if $\mathcal{D}$ is presentable and F has a fully faithful right adjoint.⁸

Proposition 2.3.7 ([47, Tag 06VH]). Let $\mathcal{C}$ be a presentable category and W be a small collection of morphisms, and $\mathcal{C}_0 \subseteq \mathcal{C}$ be the subcategory of W -local objects. Then, $\mathcal{C}_0$ is a Bousfield localization of $\mathcal{C}$.

Remark 2.3.8. In fact, $\mathcal{C}_0$ can be characterized as the universal category with morphisms in W being inverted and is thus sometimes denoted by $\mathcal{C}[W]$. That is, there is a left adjoint $L : \mathcal{C} \rightarrow \mathcal{C}[W]$ to the inclusion $\mathcal{C}[W] := \mathcal{C}_0 \subseteq \mathcal{C}$ such that, for any category $\mathcal{D}$ with colimits, precomposing with L induces an inclusion

$$\mathrm{Fun}^{\mathrm{colim}}(\mathcal{C}[W], \mathcal{D}) \hookrightarrow \mathrm{Fun}^{\mathrm{colim}}(\mathcal{C}, \mathcal{D})$$

whose image consists of $F : \mathcal{C} \rightarrow \mathcal{D}$ that sends any $w \in W$ to an isomorphism $F(w)$.

Definition 2.3.9 ([47, Tag 04KG]). Let $w : X \rightarrow Y$ be a morphism in $\mathcal{C}$. An object $C \in \mathcal{C}$ is said to be w -local if the induced map

$$(-) \circ w : \mathrm{Hom}_{\mathcal{C}}(Y, C) \xrightarrow{\sim} \mathrm{Hom}_{\mathcal{C}}(X, C)$$

is an isomorphism. For a collection of morphisms W , an object $C \in \mathcal{C}$ is said to be W -local if it is w -local for all $w \in W$.

Theorem 2.3.10. Let X be a topological space. The inclusion $\mathrm{Sh}(X) \subseteq \mathrm{PSh}(X)$ is a Bousfield localization. In particular, $\mathrm{Sh}(X)$ is presentable.

Proof. Let $U \subseteq X$ be an open set. The (strong) Yoneda Lemma, (1) of Theorem 2.1.15, implies that $F(U) = \mathrm{Hom}_{\mathrm{PSh}(X)}(h_U, F)$. Note that this embedding is covariant: if $V \subseteq U$ is an inclusion of opens, then we have the unique map $h_V \rightarrow h_U$. Now, let $\{U_\alpha\}_{\alpha \in I}$ be an open cover of U , we see that there is a diagram

$$\cdots \rightrightarrows \coprod_{\alpha, \beta, \gamma \in I} h_{U_{\alpha\beta\gamma}} \rightrightarrows \coprod_{\alpha, \beta \in I} h_{U_{\alpha\beta}} \rightrightarrows \coprod_{\alpha \in I} h_{U_\alpha} \rightarrow h_U$$

and thus the morphism

$$(2) \quad \mathrm{colim} \left(\cdots \rightrightarrows \coprod_{\alpha, \beta, \gamma \in I} h_{U_{\alpha\beta\gamma}} \rightrightarrows \coprod_{\alpha, \beta \in I} h_{U_{\alpha\beta}} \rightrightarrows \coprod_{\alpha \in I} h_{U_\alpha} \right) \rightarrow h_U$$

at which one evaluates F to get the sheaf condition (1). In other words, the subcategory of sheaves consists exactly of objects local with respect to all such morphisms. $\square$

Definition 2.3.11. By Theorem 2.3.10, the inclusion $\mathrm{Sh}(X) \hookrightarrow \mathrm{PSh}(X)$ has a left adjoint and we will refer to it as the sheafification functor. Furthermore, for a presheaf F , we will denote its sheafification by $F^\dagger$; the symbol is pronounced as “dagger”.

⁸This is called accessible localization in [44].

Remark 2.3.12. In fact, every presentable category is a Bousfield localization of a presheaf category [47, Tag 06VP]. What is special about the sheafification functor is that it has a very explicit construction [44, Remark 6.2.2.12]. This construction works more generally for sheaves on a site, and in particular shows that the left adjoint is left exact, so sheaves on a site form a topos [44, Proposition 6.2.2.7]. In our case, for a presheaf F , one can consider the assignment

$$F^\dagger(U) := \operatorname{colim}_{\mathcal{U}} \lim_{U_\alpha \in \mathcal{U}} F(U_\alpha)$$

where the filtered colimit, indexed by $\mathcal{U}$, runs over the family of all open covers $\{U_\alpha\}$ of U and, for a fixed open cover, the limit is given by (1). In the classical situation, when $\mathcal{C} = \mathbf{Set}$, this gives a description of sheafification. When $\mathcal{C} = (1,1)\mathbf{Cat}$ is the $(2,1)$ -category of classical categories, one has to do it twice and the sheafification is given by $(F^\dagger)^\dagger$. The essential content of [44, Proposition 6.2.2.7] is that, in the higher categorical setting, transfinitely iterating $(-)^\dagger$ will give the sheafification.

Our next goal is to construct the $*$ -functoriality: For a continuous map $f : X \rightarrow Y$, we will construct an adjunction $f^* : \mathbf{Sh}(Y) \rightleftarrows \mathbf{Sh}(X) : f_*$, the $*$ -pullback and $*$ -pushforward.

Construction 2.3.13. Constructing pushforward is straightforward. The map f induces a map

$$\begin{aligned} f^{-1} : \mathbf{Op}_Y &\rightarrow \mathbf{Op}_X \\ V &\mapsto f^{-1}(V) \end{aligned}$$

between the posets. Applying $\mathbf{PSh}(-)$, we obtain a pushforward

$$f_*^{\text{pre}} : \mathbf{PSh}(X) \rightarrow \mathbf{PSh}(Y)$$

on the level of presheaves and on objects it is given by the formula $(f_*^{\text{pre}} F)(V) = F(f^{-1}(V))$. The observation is that its restriction to $\mathbf{Sh}(X)$ factors through $\mathbf{Sh}(Y)$, as an open cover of V pulls back to an open cover of $f^{-1}(V)$ and the sheaf condition (1) of F implies that of $f_* F$. In summary, we have a $*$ -pushforward functor $f_* : \mathbf{Sh}(X) \rightarrow \mathbf{Sh}(Y)$.

Remark 2.3.14. Observe that f_*^{pre} is a special example of the functor obtained from precomposition discussed in Example 2.2.5. However, the notation is somehow reversed since we reverse the arrow one extra time when passing from $f : X \rightarrow Y$ to $f^{-1} : \mathbf{Op}_Y \rightarrow \mathbf{Op}_X$. Both notations are, unfortunately, standard.

Construction 2.3.15. As noted in Remark 2.3.14, as f_*^{pre} is a precomposition, it admits a left adjoint $f^{*,\text{pre}} : \mathbf{PSh}(Y) \rightarrow \mathbf{PSh}(X)$ by the left Kan extension. Recall from Example 2.2.5, on objects, it uses opens $V \subseteq Y$ to approximate opens $U \subseteq X$. (Note however that, since $\mathbf{PSh}(X) = \mathbf{Fun}(\mathbf{Op}_X^{\text{op}}, \mathbf{Ani})$, there is an extra op that flips the direction of the arrows.) More concretely, for $G \in \mathbf{PSh}(Y)$ and $U \subseteq X$, we have

$$f^{*,\text{pre}}(G)(U) = \operatorname{colim}_{f^{-1}(V) \supseteq U} G(V)$$

where V runs over open sets in Y that contains $f^{-1}(U)$. Equivalently, since $f(U)$ might not be an open set in Y , one runs over all its open neighborhoods V and assigns the anima by those $G(V)$.

Now, as f_* is restricted from f_*^{pre} , the left adjoint $f^* : \mathbf{Sh}(Y) \rightarrow \mathbf{Sh}(X)$ is given by the composition

$$\mathbf{Sh}(Y) \hookrightarrow \mathbf{PSh}(Y) \xrightarrow{f^{*,\text{pre}}} \mathbf{PSh}(X) \xrightarrow{(-)^\dagger} \mathbf{Sh}(X).$$

Example 2.3.16. Let $i : S \subseteq X$ be a subset and we endow it with the subspace topology. For a presheaf $G \in \text{PSh}(X)$, we claim that the counit $i^{*,\text{pre}} i_*^{\text{pre}} \rightarrow \text{id}$ is an isomorphism. In particular, if $G \in \text{Sh}(S)$ is a sheaf, then $i^{*,\text{pre}}(i_* G) = G$ is already a sheaf and no sheafification is needed. Recall an open set $W \subseteq S$ is a set of the form $W = U \cap S$ where $U \subseteq X$ is open. We then compute that the unit is given by

$$i^{*,\text{pre}}(i_* G)(W) = \text{colim}_{V \supseteq W} \Gamma(V; i_* G) = \text{colim}_{V \supseteq W} \Gamma(V \cap S; G) \rightarrow \Gamma(W; G).$$

But by cofinality, we can restrict the colimit to $V \subseteq U$ and it would imply that $V \cap S \subseteq W$ so the filtered diagram takes constant value and the last map is an isomorphism.

Corollary 2.3.17. *For a continuous map $f : X \rightarrow Y$, there is a $*$ -adjunction*

$$f^* : \text{Sh}(Y) \rightleftarrows \text{Sh}(X) : f_*$$

where f_* is constructed in Construction 2.3.13.

Before we leave this section, we mention a few properties of $\text{Sh}(-)$ that we will be using later. First, “sheaves form a sheaf”.

Proposition 2.3.18. *The presheaf*

$$\begin{aligned} \text{Sh}^* : \text{Op}_X^{\text{op}} &\rightarrow \text{Pr}^{\text{L}} \\ U &\mapsto \text{Sh}(U) \\ (V \subseteq U) &\mapsto (\text{Sh}(U) \rightarrow \text{Sh}(V)) \end{aligned}$$

is a sheaf in Pr^{L} .

Proof. By Theorem 2.3.10, $\text{Sh}(X)$ is a localization of $\text{PSh}(X)$ with respect to a Grothendieck topology. By [44, Proposition 6.2.2.7], any such localization is a topos and the desired statement is a consequence of (ii) of [44, Theorem 6.1.3.9]. $\square$

Another property is the base-change formula for proper maps. Consider a commuting diagram

$$\begin{array}{ccc} X' & \xrightarrow{q'} & X \\ \downarrow p' & & \downarrow p \\ Y' & \xrightarrow{q} & Y \end{array}$$

in topological spaces. The identification $p_* q'_* = q_* p'_*$ induces a map

$$(3) \quad q^* p_* \rightarrow q^* p_* q'_* q'^* = q^* q_* p'_* q'^* \rightarrow p'_* q'^*$$

between functors from $\text{Sh}(X)$ to $\text{Sh}(Y')$.

Theorem 2.3.19 ([44, Corollary 7.3.1.18], the nonabelian base change theorem).

$$\begin{array}{ccc} X' & \xrightarrow{q'} & X \\ \downarrow p' & & \downarrow p \\ Y' & \xrightarrow{q} & Y \end{array}$$

be a pullback diagram of locally compact Hausdorff spaces and assume p is proper. Then, p_* satisfies base-change, i.e., the canonical map $q^* p_* \rightarrow p'_* q'^* : \text{Sh}(X) \rightarrow \text{Sh}(Y')$, defined by (3), is an equivalence.

Remark 2.3.20. What is proven more generally in [44, Section 7.3] is that a proper map $p : X \rightarrow Y$ induces a proper morphism $p_* : \mathrm{Sh}(X) \rightarrow \mathrm{Sh}(Y)$ between topoi. In addition to base-change above, this implies for example that the $*$ -pushforward $p_* : \mathrm{Sh}(X) \rightarrow \mathrm{Sh}(Y)$ preserves filtered colimits. See [44, Remark 7.3.1.5] and [44, Corollary 7.3.4.12].

2.4. The stable setting. Classically, in order to apply sheaf theory to study geometry, instead of working with set-valued sheaves $\mathrm{Sh}(X; \mathrm{Set})$, we would usually work with abelian-group valued sheaves $\mathrm{Sh}(X; \mathrm{Ab})$. The reason is that an abelian category, having finite products agree with finite coproducts combined with the existence of well-behaved kernels and cokernels, which ensure all finite limits and colimits exist and behave well. We note that, although sometimes an abelian category $\mathcal{C}$ is defined as a category enriched over abelian groups satisfying certain properties, the enrichment assumption is in fact redundant and can be recovered from the property of being abelian. In other words, a category being abelian is a property not an extra structure.

This is in contrast with the usual triangulated categories formulation where the class of distinguished triangles is an extra structure, which is a way to get around the issues of lacking limits and colimits as illustrated in Example 2.1.1. However, this approach soon runs into the problem of not gluing well as seen in Example 2.1.2. One higher categorical solution is to define a higher categorical analogue of abelian categories, and one such notion is that of stable categories. The standard reference is [45, Chapter 1].

Definition 2.4.1. A category $\mathcal{C}$ is said to be pointed if it has a zero object, i.e., an object 0 which is both initial and final. In a pointed category, we say that a sequence $A \xrightarrow{g} B \xrightarrow{f} C$ is a fiber sequence if the commutative diagram

$$\begin{array}{ccc} A & \xrightarrow{g} & B \\ \downarrow & & \downarrow f \\ 0 & \longrightarrow & C \end{array}$$

is a pullback. We say that the same sequence $A \xrightarrow{g} B \xrightarrow{f} C$ is a cofiber sequence if the above commutative diagram is a pushout. A category $\mathcal{C}$ is said to be stable if it is pointed and if any sequence $A \xrightarrow{g} B \xrightarrow{f} C$ is a fiber sequence if and only if it is a cofiber sequence.

Lemma 2.4.2 ([45, Proposition 1.1.3.4]). *Let $\mathcal{C}$ be a pointed category. Then $\mathcal{C}$ is stable if and only if the following conditions are satisfied*

- (1) *The category $\mathcal{C}$ has finite limits and colimits.*
- (2) *Any commutative diagram*

$$\begin{array}{ccc} A & \longrightarrow & B \\ \downarrow & & \downarrow \\ D & \longrightarrow & C \end{array}$$

is a pullback if and only if it is a pushout.

*In particular, in a stable category $\mathcal{C}$, a finite limit is also a finite colimit, and finite limits and colimits thus commute with all limits and colimits.*⁹

⁹See [44, Corollary 4.4.2.5] and [44, Proposition 4.4.2.6].

Definition 2.4.3. We say a functor $F : \mathcal{C} \rightarrow \mathcal{D}$ between stable categories is exact if it sends fiber sequences to fiber sequences.

Lemma 2.4.4 ([45, Proposition 1.1.4.1]). *Let $F : \mathcal{C} \rightarrow \mathcal{D}$ be a functor between stable categories. The following are equivalent.*

- (1) F is exact.
- (2) F is left exact, i.e., it preserves finite limits.
- (3) F is right exact, i.e., it preserves finite colimits.

Notation 2.4.5. We will use the notation $\text{st} \subseteq \text{Cat}$ to denote the non-full subcategory of (idempotent complete¹⁰) stable categories with morphisms given by exact functors.

Remark 2.4.6. If $\mathcal{C}$ is a stable category, then its homotopy category $h\mathcal{C}$ has a canonical triangulated structure such that a sequence $A \rightarrow B \rightarrow C \xrightarrow{+1} A$ in $h\mathcal{C}$ is a distinguished triangle if and only if $A \rightarrow B \rightarrow C$ is a fiber sequence. There exist triangulated categories that do not come from a stable category. However, most of the natural examples do. For example in [45, 1.3.5], a construction of the (unbounded) derived category $D(\mathcal{A})$ when $\mathcal{A}$ is a Grothendieck abelian category is given. Such class includes the category of modules over a ring Mod_R and its bounded part recovers the classical notion after taking h . More related to us, if M is a manifold, the derived category $D(M; \mathbb{Z})$ is the homotopy category of $\text{Sh}(M; \text{Mod}_{\mathbb{Z}})$.¹¹

There is a general procedure to obtain a stable category from an unstable one.¹² A few equivalent constructions can be made. But we will choose one that is most geometric.

Definition 2.4.7. Let $\mathcal{C}$ be a pointed category with finite limits and colimits. For an object $X \in \mathcal{C}$, we define its loop (space) $\Omega_{\mathcal{C}}X$ by the pullback diagram.

$$\begin{array}{ccc} \Omega_{\mathcal{C}}X & \longrightarrow & 0 \\ \downarrow & & \downarrow \\ 0 & \longrightarrow & X. \end{array}$$

Equivalently, $\Omega_{\mathcal{C}}X$ is the fiber of the canonical map $0 \rightarrow X$. Similarly, we define its suspension $\Sigma_{\mathcal{C}}X$ by the cofiber sequence $X \rightarrow 0 \rightarrow \Sigma X$. Note that as these objects are defined by either a limit or a colimit, $\Omega_{\mathcal{C}}$ and $\Sigma_{\mathcal{C}}$ automatically organize into endofunctors $\Omega_{\mathcal{C}} : \mathcal{C} \rightarrow \mathcal{C}$ and $\Sigma_{\mathcal{C}} : \mathcal{C} \rightarrow \mathcal{C}$. One can also check that $\Sigma_{\mathcal{C}} \dashv \Omega_{\mathcal{C}}$ forms an adjunction pair.

Lemma 2.4.8 ([45, Proposition 1.4.2.11]). *Let $\mathcal{C}$ be a pointed category which admits finite limits and colimits. Then, $\mathcal{C}$ is stable if and only if the loop functor $\Omega_{\mathcal{C}} : \mathcal{C} \rightarrow \mathcal{C}$ is an equivalence.*

Remark 2.4.9. We note that in a stable category, there is a fiber sequence $X \rightarrow Y \rightarrow Z$ if and only if there is a fiber sequence $\Omega_{\mathcal{C}}Z \rightarrow X \rightarrow Y$ if and only if there is a fiber sequence $Y \rightarrow Z \rightarrow \Sigma_{\mathcal{C}}X$. This is because pullback and pushout diagrams satisfy the 2-out-of-3 rules

¹⁰More on this condition later.

¹¹In general, the derived category of the abelian category of $\mathbb{Z}$ -valued sheaves would give hypersheaves on M . But all sheaves satisfy hypersheaves when the space is finite dimensional [44, Chapter 7].

¹²Here we use the convention of having unstable to mean not-necessarily stable.

and one can consider for example the following commutative diagram:

$$\begin{array}{ccccc} \Omega_{\mathcal{C}} Z & \longrightarrow & X & \longrightarrow & 0 \\ \downarrow & & \downarrow & & \downarrow \\ 0 & \longrightarrow & Y & \longrightarrow & Z. \end{array}$$

Notation 2.4.10. When $\mathcal{C}$ is stable, for $X \in \mathcal{C}$, we sometimes denote $\Omega_{\mathcal{C}}(X)$ as $X[-1]$ and $\Sigma_{\mathcal{C}}(X)$ as $X[1]$. We use a similar notation $X[n]$ for $n \in \mathbb{Z}$ to denote the result of iterating these functors.

According to the lemma, given any pointed category $\mathcal{C}$, we can try to stabilize $\mathcal{C}$ by formally inverting $\Omega_{\mathcal{C}}$.

Construction 2.4.11 ([45, Proposition 1.4.2.24]). Let $\mathcal{C}$ be a pointed category with finite limits. As mentioned in Proposition 2.1.10, the category Cat admits limits and we define $\text{Sp}(\mathcal{C})$, the category of spectrum objects in $\mathcal{C}$, to be the limit

$$\text{Sp}(\mathcal{C}) := \lim \left(\cdots \xrightarrow{\Omega_{\mathcal{C}}} \mathcal{C} \xrightarrow{\Omega_{\mathcal{C}}} \mathcal{C} \xrightarrow{\Omega_{\mathcal{C}}} \mathcal{C} \right),$$

and we also use the notation $\Omega_{\mathcal{C}}^{\infty} : \text{Sp}(\mathcal{C}) \rightarrow \mathcal{C}$ to denote the projection to the last component. Per Remark 2.1.11, an object $X \in \text{Sp}(\mathcal{C})$ is a family $(X_n)_{n=0}^{\infty}$ such that $X_n = \Omega_{\mathcal{C}}(X_{n+1})$ and in this case, $\Omega_{\mathcal{C}}^{\infty}(X) = X_0$. For this reason, $\text{Sp}(\mathcal{C})$ is also called the category of spectrum objects in $\mathcal{C}$.

Remark 2.4.12. We note that by taking limits along the morphism between the diagrams

$$\begin{array}{ccccccc} \cdots & \xrightarrow{\Omega_{\mathcal{C}}} & \mathcal{C} & \xrightarrow{\Omega_{\mathcal{C}}} & \mathcal{C} & \xrightarrow{\Omega_{\mathcal{C}}} & \mathcal{C} \\ & & \downarrow \Omega_{\mathcal{C}} & & \downarrow \Omega_{\mathcal{C}} & & \downarrow \Omega_{\mathcal{C}} \\ \cdots & \xrightarrow{\Omega_{\mathcal{C}}} & \mathcal{C} & \xrightarrow{\Omega_{\mathcal{C}}} & \mathcal{C} & \xrightarrow{\Omega_{\mathcal{C}}} & \mathcal{C} \end{array}$$

that the endomorphism $\Omega_{\mathcal{C}}$ naturally extends to an endomorphism, which by abuse of notation we also denote by $\Omega_{\mathcal{C}} : \text{Sp}(\mathcal{C}) \rightarrow \text{Sp}(\mathcal{C})$. Observe also that according to the construction, for $n \geq 0$, $\Omega_{\mathcal{C}}^{\infty} \circ \Omega_{\mathcal{C}}^n = \Omega_{\mathcal{C}}^n \circ \Omega_{\mathcal{C}}^{\infty}$ and we use the notation $\Omega_{\mathcal{C}}^{\infty+n}$ to denote it. But, as X_{k+1} plays the role of delooping for X_k for an object $X = (X_k)_{k=0}^{\infty}$, we will also use the notation $\Omega_{\mathcal{C}}^{\infty-n} : \text{Sp}(\mathcal{C}) \rightarrow \mathcal{C}$ to denote the projection that on objects gives $(X \mapsto X_n)$. Now, the claim is that the induced endomorphism of the diagram

$$\begin{array}{ccccccc} \cdots & \xrightarrow{\text{id}_{\text{Sp}(\mathcal{C})}} & \text{Sp}(\mathcal{C}) & \xrightarrow{\text{id}_{\text{Sp}(\mathcal{C})}} & \text{Sp}(\mathcal{C}) & \xrightarrow{\text{id}_{\text{Sp}(\mathcal{C})}} & \text{Sp}(\mathcal{C}) \\ & & \downarrow \Omega_{\mathcal{C}}^{\infty-3} & & \downarrow \Omega_{\mathcal{C}}^{\infty-2} & & \downarrow \Omega_{\mathcal{C}}^{\infty-1} \\ \cdots & \xrightarrow{\Omega_{\mathcal{C}}} & \mathcal{C} & \xrightarrow{\Omega_{\mathcal{C}}} & \mathcal{C} & \xrightarrow{\Omega_{\mathcal{C}}} & \mathcal{C} \end{array}$$

provides the inverse to $\Omega_{\mathcal{C}} : \text{Sp}(\mathcal{C}) \rightarrow \text{Sp}(\mathcal{C})$. More concretely, the effect of $\Omega_{\mathcal{C}}$ on objects is given by

$$(\cdots, X_2, X_1, X_0) \mapsto (\cdots, \Omega_{\mathcal{C}} X_2, \Omega_{\mathcal{C}} X_1, \Omega_{\mathcal{C}} X_0) = (\cdots, X_1, X_0, \Omega_{\mathcal{C}} X_0)$$

and thus has an inverse given by shifting upward

$$(\cdots, Y_2, Y_1, Y_0) \mapsto (\cdots, Y_3, Y_2, Y_1).$$

Example 2.4.13. In the case when $\mathcal{C} = \text{Ani}_*$ is the category of pointed anima, the resulting category is usually denoted simply as Sp and is referred to as the category of spectra. This category plays an important role in classical stable homotopy theory. Indeed, an object $X = (X_n)_{n=0}^\infty \in \text{Sp}$ satisfies $X_0 = \Omega_*(X_1) = \Omega_*^2(X_2) = \cdots$ so it gives an infinite loop space. See the introduction of [45, Section 1.4] or [45, 1.4.3] for more details.

Here, we only mention that the fact that $\pi_0(\Omega_*^n X) = \pi_n(X)$ for $X \in \text{Ani}_*$ implies that, for $X \in \text{Sp}$, there is a notion of homotopy group $\pi_n(X)$ for $n \in \mathbb{Z}$, even the negative numbers. Indeed, with the expression $X = (X_n)$, for $m \in \mathbb{N}$, $\pi_{-m}(X) := \pi_0(X_m) = \pi_1(X_{m+1}) = \cdots$.

Remark 2.4.14. Recall that an abelian category $\mathcal{A}$ is naturally enriched in the category of abelian groups Ab . A similar situation holds for stable categories: By Lemma 2.4.8, if $\mathcal{C}$ is a stable category, then the functors $\Sigma_{\mathcal{C}}$ and $\Omega_{\mathcal{C}}$, which form an adjoint pair, are in fact inverse to each other. Thus, for $X, Y \in \mathcal{C}$, the sequence

$$(\cdots, \text{Hom}_{\mathcal{C}}(X, \Sigma_{\mathcal{C}}^n(Y)), \cdots, \text{Hom}_{\mathcal{C}}(X, \Sigma_{\mathcal{C}}(Y)), \text{Hom}_{\mathcal{C}}(X, Y))$$

forms a spectrum in Sp . For example, as $\Omega_{\mathcal{C}}$ is defined as a limit, we have

$$\Omega_*(\text{Hom}_{\mathcal{C}}(X, \Sigma_{\mathcal{C}}(Y))) = \text{Hom}_{\mathcal{C}}(X, \Omega_{\mathcal{C}}\Sigma_{\mathcal{C}}(Y)) = \text{Hom}_{\mathcal{C}}(X, Y).$$

Lemma 2.4.15. *Let $\mathcal{C}$ be a pointed category with limits. The category $\text{Sp}(\mathcal{C})$ is stable.*

Proof. The proof is taken from [45, Proposition 1.4.2.24]. One first notices that $\text{Sp}(\mathcal{C})$ is pointed and, per Remark 2.4.12, $\Omega_{\mathcal{C}} : \text{Sp}(\mathcal{C}) \rightarrow \text{Sp}(\mathcal{C})$ is invertible. However, we note that Lemma 2.4.8 cannot be applied directly, since we do not assume $\mathcal{C}$ to have colimits, but we can consider the following maneuver.

Assume the case that $\mathcal{C}$ is presentable. As remarked in Definition 2.4.7, the loop functor $\Omega_{\mathcal{C}}$ is a right adjoint. Thus, the limit diagram in Construction 2.4.11 in fact can be taken in Pr^{R} so $\text{Sp}(\mathcal{C})$ in this case is presentable and has both limits and colimits by Proposition 2.2.19. In other words, Lemma 2.4.8 applies when $\mathcal{C}$ is presentable and we know that $\text{Sp}(\mathcal{C})$ is stable.

For the general case, consider the Yoneda embedding $j : \mathcal{C} \hookrightarrow \text{PSh}(\mathcal{C})$, which we know is fully faithful and left exact. By construction, we have $\text{Sp}(j) \circ \Omega_{\mathcal{C}} = \Omega_{\text{PSh}(\mathcal{C})} \circ \text{Sp}(j)$. Additionally, since hom-anima in a limit category can be computed explicitly by Remark 2.1.11, we conclude that $\text{Sp}(j) : \text{Sp}(\mathcal{C}) \hookrightarrow \text{Sp}(\text{PSh}(\mathcal{C}))$ is also fully faithful and left exact. But this implies that finite colimits computed in $\text{Sp}(\text{PSh}(\mathcal{C}))$ in fact exist in $\mathcal{C}$ already. For example, for a morphism $f : X \rightarrow Y$ in $\mathcal{C}$, the cofiber is equivalent to $\Omega_{\mathcal{C}}^{-1}(\text{fib}(f))$, as one can check its universal property after embedding it to $\text{Sp}(\text{PSh}(\mathcal{C}))$. $\square$

Recall that, for any category $\mathcal{C}$ with finite limits, the category of pointed objects $\mathcal{C}_*$ in $\mathcal{C}$ is the universal pointed category that maps to $\mathcal{C}$. In particular, if $\mathcal{C}$ is already pointed, then the forgetful map $\mathcal{C}_* \rightarrow \mathcal{C}$ is an equivalence.

Definition 2.4.16. For a category $\mathcal{C}$ with finite limits, we define its stabilization $\text{Sp}(\mathcal{C})$ to be $\text{Sp}(\mathcal{C}_*)$ where the latter is given by Construction 2.4.11.¹³ Furthermore, we also use the notation $\Omega_{\mathcal{C}}^\infty : \text{Sp}(\mathcal{C}) \rightarrow \mathcal{C}$ for the composition

$$\text{Sp}(\mathcal{C}) := \text{Sp}(\mathcal{C}_*) \xrightarrow{\Omega_{\mathcal{C}_*}^\infty} \mathcal{C}_* \rightarrow \mathcal{C}.$$

¹³See [45, Definition 1.4.2.8] for the official definition, as reduced excisive functors out of finite spaces, and [45, Proposition 1.4.2.24] for the comparison.

Proposition 2.4.17. *For a category $\mathcal{C}$ with finite limits, its stabilization $\mathrm{Sp}(\mathcal{C})$ is universal among stable categories. That is, for any stable category $\mathcal{D}$, composing with $\Omega_{\mathcal{C}}^{\infty} : \mathrm{Sp}(\mathcal{C}) \rightarrow \mathcal{C}$ induces an equivalence*

$$\mathrm{Fun}^{\mathrm{lex}}(\mathcal{D}, \mathrm{Sp}(\mathcal{C})) \xrightarrow{\sim} \mathrm{Fun}^{\mathrm{lex}}(\mathcal{D}, \mathcal{C})$$

where $\mathrm{Fun}^{\mathrm{lex}}(\mathcal{D}, \mathcal{C}) \subseteq \mathrm{Fun}(\mathcal{D}, \mathcal{C})$ is the subcategory consisting of left exact functors.

We have been restricting our concern to only stability, but let us now consider its interaction with presentability. We first notice that within the stable setting, checking whether a functor preserves colimits (and limits) is easier. In fact, we have

Proposition 2.4.18 ([45, Proposition 1.4.4.1]). *Let $\mathcal{C}$ and $\mathcal{D}$ be stable categories and $F : \mathcal{C} \rightarrow \mathcal{D}$ be an exact functor.*

- (1) *The category $\mathcal{C}$ admits small colimits if and only if it admits small coproducts.*
- (2) *Assume $\mathcal{C}$ and $\mathcal{D}$ admit small colimits. The functor F preserves small colimits if and only if it preserves small coproducts.*
- (3) *Let $X \in \mathcal{C}$ be an object. Then X is compact if and only if the following condition holds: For any map $f : X \rightarrow \coprod_{\alpha \in A} Y_{\alpha}$, there exists a finite set A_0 such that f factorizes to $f_0 : X \rightarrow \coprod_{\alpha \in A_0} Y_{\alpha}$.*

Corollary 2.4.19. *Let $\mathcal{C}$ and $\mathcal{D}$ be stable categories and $F : \mathcal{C} \rightarrow \mathcal{D}$ be an exact functor.*

- (1) *The category $\mathcal{C}$ admits small colimits if and only if it admits filtered colimits.*
- (2) *Assume $\mathcal{C}$ and $\mathcal{D}$ admit small colimits. The functor F preserves small colimits if and only if it preserves filtered colimits.*

Proof. We prove (1). It is enough to show that $\mathcal{C}$ admits coproducts, so we consider a family $\{X_{\alpha}\}_{\alpha \in I}$, indexed by a (small) set I . Denote by $S := \{I_0 | I_0 \subseteq I, |I_0| < \infty\}$, the poset consisting of finite subsets of I with the partial order given by inclusion. Then one notices that S is filtered, the filtered colimit $\mathrm{colim}_{I_0 \in S} (\coprod_{\alpha \in I_0} X_{\alpha}) \xrightarrow{\sim} \coprod_{\alpha \in I} X_{\alpha}$ computes the desired coproduct. $\square$

The stabilization functor also works well with presentability.

Proposition 2.4.20 ([45, Proposition 1.4.4.4]). *Let $\mathcal{C}$ and $\mathcal{D}$ be presentable and assume $\mathcal{D}$ is stable. Then*

- (1) *The category $\mathrm{Sp}(\mathcal{C})$ is presentable.*
- (2) *The loop functor $\Omega_{\mathcal{C}}^{\infty} : \mathrm{Sp}(\mathcal{C}) \rightarrow \mathcal{C}$ admits a left adjoint $\Sigma_{\mathcal{C},+}^{\infty} : \mathcal{C} \rightarrow \mathrm{Sp}(\mathcal{C})$.*
- (3) *An exact functor $G : \mathcal{D} \rightarrow \mathrm{Sp}(\mathcal{C})$ admits a left adjoint if and only if $\Omega_{\mathcal{C}}^{\infty} \circ G : \mathcal{D} \rightarrow \mathcal{C}$ admits a left adjoint.*

In particular, the infinite suspension functor $\Sigma_{\mathcal{C},+}^{\infty} : \mathcal{C} \rightarrow \mathrm{Sp}(\mathcal{C})$ is universal among stable presentable categories. That is, precomposition with $\Sigma_{\mathcal{C},+}^{\infty}$ induces an equivalence

$$\mathrm{Fun}^L(\mathrm{Sp}(\mathcal{C}), \mathcal{D}) \xrightarrow{\sim} \mathrm{Fun}^L(\mathcal{C}, \mathcal{D})$$

where $\mathrm{Fun}^L(\mathcal{C}, \mathcal{D}) \subseteq \mathrm{Fun}(\mathcal{C}, \mathcal{D})$ is the subcategory consisting of colimit-preserving functors.

Notation 2.4.21. We will use the notation $\mathrm{Pr}_{st}^L \subseteq \mathrm{Pr}^L$ to denote the full subcategory consisting of stable presentable categories.

Example 2.4.22. Let X be a topological space. We claim that $\mathrm{Sp}(\mathrm{Sh}(X)) = \mathrm{Sh}(X; \mathrm{Sp})$. We begin with the level of presheaves. In fact, the situation holds more generally for any category $\mathcal{C}$ with finite limits. First, on the level of presheaves, the canonical map $\mathrm{PSh}(X; \mathcal{C})_* \rightarrow \mathrm{PSh}(X; \mathcal{C}_*)$ is an equivalence, as the final object in $\mathrm{PSh}(X; \mathcal{C})$ is given by the constant presheaf ($U \mapsto *$). So we assume $\mathcal{C}$ is pointed. But then, as $\mathrm{Fun}(-, -)$ is the internal hom in Cat , $\mathrm{PSh}(X; -) = \mathrm{Fun}(\mathrm{Op}_X^{\mathrm{op}}, -)$ preserves limits so

$$\begin{aligned} \mathrm{PSh}(X; \mathrm{Sp}(\mathcal{C})) &= \lim \left(\cdots \xrightarrow{\Omega_{\mathcal{C}}} \mathrm{Fun}(\mathrm{Op}_X^{\mathrm{op}}, \mathcal{C}) \xrightarrow{\Omega_{\mathcal{C}}} \mathrm{Fun}(\mathrm{Op}_X^{\mathrm{op}}, \mathcal{C}) \xrightarrow{\Omega_{\mathcal{C}}} \mathrm{Fun}(\mathrm{Op}_X^{\mathrm{op}}, \mathcal{C}) \right) \\ &= \mathrm{Sp}(\mathrm{PSh}(X; \mathcal{C})). \end{aligned}$$

Concretely, a sequence of assignments $U \mapsto F_n(U)$ satisfying $\Omega F_{n+1}(U) = F_n(U)$, i.e., a spectrum object in presheaf, corresponds exactly to a presheaf of spectrum objects. Under this identification, as sheaves are presheaves with properties given by limits with a special form, we see that the inclusion $\mathrm{Sh}(X; \mathcal{C}) \hookrightarrow \mathrm{PSh}(X; \mathcal{C})$ stabilizes to $\mathrm{Sh}(X; \mathrm{Sp}(\mathcal{C})) \hookrightarrow \mathrm{PSh}(X; \mathrm{Sp}(\mathcal{C}))$.

We can now stabilize to obtain the following:

Corollary 2.4.23. *The category $\mathrm{Sh}(X; \mathrm{Sp})$ is presentable. The sheafification functor, the left adjoint of the inclusion $\mathrm{Sh}(X; \mathrm{Sp}) \hookrightarrow \mathrm{PSh}(X; \mathrm{Sp})$, exists for spectra-valued sheaves. For a continuous map $f : X \rightarrow Y$, there is a $*$ -adjunction*

$$f^* : \mathrm{Sh}(Y; \mathrm{Sp}) \rightleftarrows \mathrm{Sh}(X; \mathrm{Sp}) : f_*$$

where f_* is constructed the same way as in Construction 2.3.13.

2.5. Künneth and duality. The next general machinery we need in order to develop the six-functor formalism for sheaves in topology is the notion of tensor product for categories.

Definition 2.5.1. We will use the notation $\widehat{\mathrm{Cat}}^{\mathrm{colim}}$ to denote the (very large) category consisting of (large) categories with small colimits with morphisms given by colimit-preserving functors.

As in the case of Cat , if $\mathcal{C}, \mathcal{D} \in \widehat{\mathrm{Cat}}$, then the Cartesian product $\mathcal{C} \times \mathcal{D}$ will also have small colimits.¹⁴ Indeed, Remark 2.1.11 implies that a diagram $I \rightarrow \mathcal{C} \times \mathcal{D}$ is simply the same as a pair of diagrams indexed by I in $\mathcal{C}$ and $\mathcal{D}$, say x_α and y_α , $\alpha \in I$. The same remark then implies that the object $(\mathrm{colim}_\alpha x_\alpha, \mathrm{colim}_\beta y_\beta)$ gives the colimit as

$$\begin{aligned} \mathrm{Hom}_{\mathcal{C} \times \mathcal{D}} \left(\mathrm{colim}_{\alpha \in I} (x_\alpha, y_\alpha), (u, v) \right) &:= \lim_{\alpha \in I} (\mathrm{Hom}_{\mathcal{C} \times \mathcal{D}}((x_\alpha, y_\alpha), (u, v))) \\ &= \lim_{\alpha \in I} (\mathrm{Hom}_{\mathcal{C}}(x_\alpha, u) \times \mathrm{Hom}_{\mathcal{D}}(y_\alpha, v)) = \left(\lim_{\alpha \in I} \mathrm{Hom}_{\mathcal{C}}(x_\alpha, u) \right) \times \left(\lim_{\beta \in I} \mathrm{Hom}_{\mathcal{D}}(y_\beta, v) \right) \\ &= \mathrm{Hom}_{\mathcal{C}}(\mathrm{colim}_\alpha x_\alpha, u) \times \mathrm{Hom}_{\mathcal{D}}(\mathrm{colim}_\beta y_\beta, v) = \mathrm{Hom}_{\mathcal{C} \times \mathcal{D}} \left((\mathrm{colim}_{\alpha \in I} x_\alpha, \mathrm{colim}_{\beta \in I} y_\beta), (u, v) \right). \end{aligned}$$

Here, the second equality on the second line is the statement that “limits commute with limits”. However, the category $\mathcal{C} \times \mathcal{D}$ is not the “optimal” object in $\widehat{\mathrm{Cat}}^{\mathrm{colim}}$ as it does not satisfy bilinearity: pairs of colimit-preserving functors $\mathcal{C} \rightarrow \mathcal{E}$, $\mathcal{D} \rightarrow \mathcal{E}$ do not correspond to

¹⁴We thank Yuxin Chen for pointing out a mistake related to this point in an earlier draft.

colimit-preserving functors $\mathcal{C} \times \mathcal{D} \rightarrow \mathcal{E}$. The fix is a form of completion analogous to the map $M \times N \rightarrow M \otimes N$ in module theory.

Proposition 2.5.2 ([45, Section 4.8.1]). *For $\mathcal{C}$ and $\mathcal{D}$ in $\widehat{\text{Cat}}^{\text{colim}}$, there is a category with small colimits $\mathcal{C} \otimes \mathcal{D} \in \widehat{\text{Cat}}^{\text{colim}}$ with a canonical map $\mathcal{C} \times \mathcal{D} \rightarrow \mathcal{C} \otimes \mathcal{D}$ such that, for any $\mathcal{E} \in \widehat{\text{Cat}}^{\text{colim}}$, the precomposition*

$$\text{Fun}^L(\mathcal{C} \otimes \mathcal{D}, \mathcal{E}) \hookrightarrow \text{Fun}(\mathcal{C} \times \mathcal{D}, \mathcal{E})$$

is an inclusion whose image is given by functors which are colimit-preserving on each component. Furthermore, the construction $(\mathcal{C}, \mathcal{D}) \mapsto \mathcal{C} \otimes \mathcal{D}$ organizes to a symmetric monoidal structure on $\widehat{\text{Cat}}^{\text{colim}}$.

Example 2.5.3. Consider the free functor $\text{Set} \rightarrow \text{Ab}$ that sends a set X to the free abelian group $\mathbb{Z}^{\oplus X}$ generated by X . It induces naturally a functor

$$\begin{aligned} \text{Set} \times \text{Set} &\rightarrow \text{Ab} \\ (X, Y) &\mapsto \mathbb{Z}^{\oplus(X \times Y)}. \end{aligned}$$

The functor is colimit-preserving on each component, since for a fixed set X , taking product with $X \times (-) : \text{Set} \rightarrow \text{Set}$ preserves colimits. (It is sufficient to check the case of small coproducts and coequalizers, which are both straightforward.) However, the functor is not colimit-preserving. Take for example the colimit $(\{a\}, \{b\}) \amalg (\{c\}, \{d\}) = (\{a, c\}, \{b, d\})$, whose image is isomorphic to $\mathbb{Z}^4$. However, the coproduct, which is just the direct sum, of the separate images is $\mathbb{Z} \oplus \mathbb{Z} = \mathbb{Z}^2$. We note that, in this case, $\text{Set} \times \text{Set} \rightarrow \text{Set} \otimes \text{Set}$ is simply given by

$$\begin{aligned} \text{Set} \times \text{Set} &\rightarrow \text{Set} \\ X \times Y &\mapsto X \times Y \end{aligned}$$

and the induced map is again just the free functor $\text{Set} \rightarrow \text{Ab}$.

Notation 2.5.4. The above symmetric monoidal structure is often referred to as the Lurie tensor product.

Proposition 2.5.5 ([45, Section 4.8.1]). *We have the following facts regarding the Lurie tensor product and the (large) category of presentable categories.*

- (1) *The category of presentable categories $\text{Pr}^L \subseteq \widehat{\text{Cat}}^{\text{colim}}$ is stable under $\otimes$. Thus, there is a symmetric monoidal structure $(\text{Pr}^L, \otimes)$, whose tensor unit is given by Ani .*
- (2) *For $\mathcal{C}, \mathcal{D} \in \text{Pr}^L$, there is a canonical equivalence*

$$\mathcal{C} \otimes \mathcal{D} = \text{Fun}^R(\mathcal{C}^{\text{op}}, \mathcal{D})$$

where $\text{Fun}^R(\mathcal{C}^{\text{op}}, \mathcal{D})$ is the subcategory of functors consisting of right adjoints.

- (3) *The symmetric monoidal category $(\text{Pr}^L, \otimes)$ is closed. More precisely, for $\mathcal{C}, \mathcal{D}$, and $\mathcal{E} \in \text{Pr}^L$, we have*

$$\text{Hom}_{\text{Pr}^L}(\mathcal{C} \otimes \mathcal{D}, \mathcal{E}) = \text{Hom}_{\text{Pr}^L}(\mathcal{C}, \text{Fun}^L(\mathcal{D}, \mathcal{E})).$$

As a result, $\otimes : \text{Pr}^L \times \text{Pr}^L \rightarrow \text{Pr}^L$ preserves small colimits on each component.

- (4) The map $\Sigma_+^\infty : \mathbf{Ani} \rightarrow \mathbf{Sp}$ exhibits $\mathbf{Sp}$ as an idempotent algebra in $\mathbf{Pr}^L$. Furthermore, $\mathbf{Mod}_{\mathbf{Sp}}(\mathbf{Pr}^L) = \mathbf{Pr}_{st}^L$. Concretely, $\mathbf{Sp} \otimes \mathbf{Sp} = \mathbf{Sp}$ and, for any presentable category $\mathcal{C}$, we have

$$\mathcal{C} \otimes \mathbf{Sp} = \mathbf{Sp}(\mathcal{C}).$$

Remark 2.5.6. It might be puzzling why the expression $\mathbf{Fun}^R(\mathcal{C}^{\mathrm{op}}, \mathcal{D})$ is covariant in $\mathcal{C}$, which we know has to be the case as $\mathcal{C} \otimes \mathcal{D}$ is. As the proof presented in [45, Proposition 4.8.1.17] uses Yoneda, i.e., studying the category $\mathbf{Fun}^{bi-L}(\mathcal{C} \times \mathcal{D}, \mathcal{E})$ of bi-colimit-preserving functors from $\mathcal{C} \times \mathcal{D}$ to any given presentable category $\mathcal{E}$, it boils down to convincing oneself that the relevant expressions are contravariant.

Before tracing the identifications, we remark that, in this proof, exchanging left and right is given by two identifications: Let $\mathcal{C}_1$ and $\mathcal{C}_2$ be presentable. Then, one has the tautological identification $\mathbf{Fun}^R(\mathcal{C}_1, \mathcal{C}_2)^{\mathrm{op}} = \mathbf{Fun}^L(\mathcal{C}_1^{\mathrm{op}}, \mathcal{C}_2^{\mathrm{op}})$ by taking the opposite category and the interesting one $\mathbf{Fun}^R(\mathcal{C}_1, \mathcal{C}_2)^{\mathrm{op}} = \mathbf{Fun}^L(\mathcal{C}_2, \mathcal{C}_1)$ by passing to the right adjoint. Here, the first ‘op’ comes from the fact that, if one reverses the arrows in the categories, then the direction of natural transformation $T(c) : F_1(c) \rightarrow F_2(c)$ will also be reversed. As for passing to adjoint, we note that if $T : F_1 \rightarrow F_2$ is a natural transformation, then the induced natural transformation is given by

$$F_2^R \xrightarrow{\eta_1} F_1^R F_1 F_2^R \xrightarrow{T} F_1^R F_2 F_2^R \xrightarrow{\epsilon_2} F_1^R.$$

Now, it is not hard to see that the penultimate expression $\mathbf{Fun}^R(\mathcal{E}, \mathbf{Fun}^R(\mathcal{C}^{\mathrm{op}}, \mathcal{D}))^{\mathrm{op}}$ is contravariant in $\mathcal{C}$; one can trace the identification and conclude that it comes from applying $(\mathbf{Fun}^R(\mathcal{E}, -))^{\mathrm{op}}$ to the expression $\mathbf{Fun}^R(\mathcal{C}^{\mathrm{op}}, \mathcal{D})$, which is contravariant. More precisely, if $F : \mathcal{C} \rightarrow \mathcal{C}'$ is a left adjoint, then $F^{\mathrm{op}} : \mathcal{C}^{\mathrm{op}} \rightarrow \mathcal{C}'^{\mathrm{op}}$ is a right adjoint and we thus have the map

$$(-) \circ F^{\mathrm{op}} : \mathbf{Fun}^R(\mathcal{C}'^{\mathrm{op}}, \mathcal{D}) \rightarrow \mathbf{Fun}^R(\mathcal{C}^{\mathrm{op}}, \mathcal{D})$$

by pre-composition. The confusing step is the next one, i.e., passing to the left adjoint

$$\mathbf{Fun}^R(\mathcal{E}, \mathbf{Fun}^R(\mathcal{C}^{\mathrm{op}}, \mathcal{D}))^{\mathrm{op}} = \mathbf{Fun}^L(\mathbf{Fun}^R(\mathcal{C}^{\mathrm{op}}, \mathcal{D}), \mathcal{E}).$$

Abstractly speaking, we have the situation of presentable categories $\mathcal{A}, \mathcal{B}_1, \mathcal{B}_2$ and a right adjoint $G : \mathcal{B}_1 \rightarrow \mathcal{B}_2$, and we are asking what functor the post-composition

$$G \circ (-) : \mathbf{Fun}^R(\mathcal{A}, \mathcal{B}_1)^{\mathrm{op}} \rightarrow \mathbf{Fun}^R(\mathcal{A}, \mathcal{B}_2)^{\mathrm{op}}$$

corresponds to under $\mathbf{Fun}^R(\mathcal{A}, \mathcal{B})^{\mathrm{op}} = \mathbf{Fun}^L(\mathcal{B}, \mathcal{A})$. The answer is simple: If $\lambda : \mathcal{A} \rightarrow \mathcal{B}_1$ is a right adjoint, then $G \circ \lambda : \mathcal{A} \rightarrow \mathcal{B}_2$ is a right adjoint with its left adjoint being $(G \circ \lambda)^L = \lambda^L \circ G^L$. That is, the corresponding map is the pre-composition

$$(-) \circ G^L : \mathbf{Fun}^L(\mathcal{B}_1, \mathcal{A}) \rightarrow \mathbf{Fun}^L(\mathcal{B}_2, \mathcal{A}).$$

Remark 2.5.7. In fact, since $\mathbf{Cat}$ already has an internal hom, i.e., for any categories $\mathcal{C}, \mathcal{D}$, and $\mathcal{E}$, we have

$$\begin{aligned} \mathbf{Fun}(\mathcal{C} \times \mathcal{D}, \mathcal{E}) &= \mathbf{Fun}(\mathcal{C}, \mathbf{Fun}(\mathcal{D}, \mathcal{E})) \\ F &\mapsto (X \mapsto F(X, -)), \end{aligned}$$

the category $\widehat{\mathbf{Cat}}^{\mathrm{colim}}$ has internal hom given by $\mathbf{Fun}^{\mathrm{colim}}(\mathcal{C}, \mathcal{D})$ for $\mathcal{C}, \mathcal{D} \in \widehat{\mathbf{Cat}}^{\mathrm{colim}}$. This is already used implicitly in the proof of [45, Proposition 4.8.1.15] and the harder fact shown in [44, Proposition 5.5.3.8] is that when $\mathcal{C}, \mathcal{D} \in \mathbf{Pr}^L$, $\mathbf{Fun}^L(\mathcal{C}, \mathcal{D}) = \mathbf{Fun}^{\mathrm{colim}}(\mathcal{C}, \mathcal{D})$ is in fact presentable.

One fact about the Lurie tensor product that is relevant to us is that the functor $\mathrm{Sh}(-; \mathrm{Ani})$ respects products of (nice) topological spaces: Let X and Y be two topological spaces, the assignment

$$\begin{aligned} \mathrm{Op}_X \times \mathrm{Op}_Y &\rightarrow \mathrm{Sh}(X \times Y) \\ (U, V) &\mapsto h_{U \times V} \end{aligned}$$

extends to a map

$$\mathrm{PSh}(X) \times \mathrm{PSh}(Y) \rightarrow \mathrm{Sh}(X \times Y).$$

As it sends covers to covers in both components, we obtain an exterior product

$$(4) \quad \begin{aligned} \mathrm{Sh}(X) \times \mathrm{Sh}(Y) &\rightarrow \mathrm{Sh}(X \times Y) \\ (F, G) &\mapsto F \boxtimes G. \end{aligned}$$

Proposition 2.5.8. *Assume X is locally compact. The map (4) induces an equivalence*

$$\mathrm{Sh}(X) \otimes \mathrm{Sh}(Y) \xrightarrow{\sim} \mathrm{Sh}(X \times Y).$$

Furthermore, if $f : X' \rightarrow X$ and $g : Y' \rightarrow Y$ are two continuous maps, then the following diagram commutes

$$\begin{array}{ccc} \mathrm{Sh}(X) \otimes \mathrm{Sh}(Y) & \longrightarrow & \mathrm{Sh}(X \times Y) \\ \downarrow f^* \otimes g^* & & \downarrow (f \times g)^* \\ \mathrm{Sh}(X') \otimes \mathrm{Sh}(Y') & \longrightarrow & \mathrm{Sh}(X' \times Y'). \end{array}$$

Proof. By [44, Theorem 7.3.3.9] and [44, Proposition 7.3.1.11], there is an equivalence

$$\mathrm{Sh}(X \times Y) = \mathrm{Sh}(X; \mathrm{Sh}(Y))$$

between the anima-valued sheaves on $X \times Y$ and $\mathrm{Sh}(Y)$ -valued sheaves on X . By (2) of Proposition 2.5.5 and Lemma 2.5.11 below, for any presentable $\mathcal{C}$, we have

$$\mathrm{Sh}(X; \mathcal{C}) = \mathrm{Fun}^R(\mathrm{Sh}(X)^{\mathrm{op}}, \mathcal{C}) = \mathrm{Sh}(X) \otimes \mathcal{C}$$

so the statement is implied by the case when $\mathcal{C} = \mathrm{Sh}(Y)$. For the commutativity statement, note that by the universal property of mapping out, we can first reduce the domain to $\mathrm{Sh}(X) \times \mathrm{Sh}(Y)$ and then further reduce it to $\mathrm{Op}_X \times \mathrm{Op}_Y$, i.e., it is sufficient to show that the diagram

$$\begin{array}{ccc} \mathrm{Op}_X \times \mathrm{Op}_Y & \longrightarrow & \mathrm{Sh}(X \times Y) \\ \downarrow f^{-1} \times g^{-1} & & \downarrow (f \times g)^* \\ \mathrm{Op}_{X'} \times \mathrm{Op}_{Y'} & \longrightarrow & \mathrm{Sh}(X' \times Y'). \end{array}$$

and this is simply the (strict) equation $(f \times g)^{-1}(U \times V) = f^{-1}(U) \times g^{-1}(V)$. $\square$

We are now finally ready to construct the six-functor formalism for sheaves in topology. We will follow Volpe's treatment in [66] by utilizing a duality between sheaves and cosheaves. However, we will discuss it within the abstract framework considered in [30]. To begin with, we define the notion of cosheaf.

Definition 2.5.9. For a topological space X and a category $\mathcal{C}$, a $\mathcal{C}$ -valued precosheaf is a functor $F : \text{Op}_X \rightarrow \mathcal{C}$. A precosheaf F is a cosheaf if it satisfies the following dual version of the *local-to-global* property: For any open $U \subseteq X$ and any open cover $\{U_\alpha\}_{\alpha \in I}$, the canonical map

$$(5) \quad \text{colim} \left(\cdots \rightrightarrows \coprod_{\alpha, \beta, \gamma \in I} F(U_\alpha \cap U_\beta \cap U_\gamma) \rightrightarrows \coprod_{\alpha, \beta \in I} F(U_\alpha \cap U_\beta) \rightrightarrows \coprod_{\alpha \in I} F(U_\alpha) \right) \xrightarrow{\sim} F(U)$$

is an equivalence. As in the case of sheaves, we sometimes would simply write it as $F(U) \xleftarrow{\sim} \text{colim}_{\alpha \in I} F(U_\alpha)$. We denote by $\text{CSh}(X; \mathcal{C})$ the subcategory of cosheaves. Similarly, when $\mathcal{C} = \text{Ani}$, we will use the notation $\text{CSh}(X) := \text{CSh}(X; \text{Ani})$ for simplicity. Alternatively, one can define $\text{CSh}(X; \mathcal{C}) := (\text{Sh}(X; \mathcal{C}^{\text{op}}))^{\text{op}}$.

We will need the following generalized version of (i) of Theorem 2.2.6 for the next lemma.

Proposition 2.5.10 ([47, Tag 06QB]). *Let $\mathcal{C}$ be a presentable category and $\mathcal{D}$ be a category. Then, a functor $F : \mathcal{C} \rightarrow \mathcal{D}$ is a left adjoint if and only if F preserves colimits.*

For a topological space X , the composition $\text{Op}_X \hookrightarrow \text{PSh}(X) \xrightarrow{(-)^\dagger} \text{Sh}(X)$ induces a map

$$(6) \quad \text{Fun}^{\text{colim}}(\text{Sh}(X), \mathcal{C}) \hookrightarrow \text{Fun}^{\text{colim}}(\text{PSh}(X), \mathcal{C}) = \text{Fun}(\text{Op}_X, \mathcal{C})$$

Lemma 2.5.11 ([66, Lemma 2.22]). *Let X be a topological space and $\mathcal{C}$ be a category admitting colimits. Then, the map (6) restricts to an equivalence*

$$\text{CSh}(X; \mathcal{C}) \xrightarrow{\sim} \text{Fun}^L(\text{Sh}(X), \mathcal{C}).$$

Equivalently, a precosheaf $F : \text{Op}_X \rightarrow \mathcal{C}$ is a cosheaf if and only if its extension to the free completion $\text{PSh}(X) \rightarrow \mathcal{C}$ factorizes through $(-)^{\dagger}$. Similarly, if $\mathcal{C}$ admits limits, then there is an equivalence

$$\text{Sh}(X; \mathcal{C}) \xrightarrow{\sim} \text{Fun}^R(\text{Sh}(X)^{\text{op}}, \mathcal{C}).$$

Proof. Note that the statement regarding $\mathcal{C}$ -valued sheaves can be obtained from the first statement by replacing $\mathcal{C}$ with $\mathcal{C}^{\text{op}}$, and notice that, for any categories $\mathcal{C}$ and $\mathcal{D}$, we have

$$(\text{Fun}^L(\mathcal{C}, \mathcal{D}))^{\text{op}} = \text{Fun}^R(\mathcal{C}^{\text{op}}, \mathcal{D}^{\text{op}})$$

so we prove the statement regarding $\mathcal{C}$ -valued cosheaves. We first compute that $F \in \text{Fun}^{\text{colim}}(\text{Sh}(X), \mathcal{C})$, when composed with (6), gives the precosheaf $\bar{F}(U) := F(h_U)$. Then we recall, by Remark 2.3.8 and (2), that as the category of sheaves is a localization,

$$\text{Fun}^{\text{colim}}(\text{Sh}(X), \mathcal{C}) \subseteq \text{Fun}^{\text{colim}}(\text{PSh}(X), \mathcal{C}) = \text{Fun}(\text{Op}_X, \mathcal{C})$$

are exactly given by F that satisfies $\text{colim}_{\alpha} F(h_{U_\alpha}) \xrightarrow{\sim} F(h_U)$, and this is exactly the cosheaf condition $\text{colim}_{\alpha} \bar{F}(U_\alpha) \xrightarrow{\sim} \bar{F}(U)$. Lastly, we use Proposition 2.5.10 to obtain that the inclusion

$$\text{Fun}^L(\text{Sh}(X), \mathcal{C}) \subseteq \text{Fun}^{\text{colim}}(\text{Sh}(X), \mathcal{C})$$

is an equivalence. □

We take a brief detour to discuss dualizability:

Definition 2.5.12. In a symmetric monoidal category $(\mathcal{C}, \otimes)$, an object $X \in \mathcal{C}$ is said to be dualizable if the functor $X \otimes (-)$ has a right adjoint of the form $Y \otimes (-)$ for some $Y \in \mathcal{C}$.

In the fashion of Remark 2.2.4, this means that there is a unit $\eta : 1_{\mathcal{C}} \rightarrow X \otimes Y$ and a counit $\epsilon : Y \otimes X \rightarrow 1_{\mathcal{C}}$ such that the compositions

$$X \xrightarrow{\eta \otimes \text{id}_X} X \otimes Y \otimes X \xrightarrow{\text{id}_X \otimes \epsilon} X$$

and

$$Y \xrightarrow{\text{id}_Y \otimes \eta} Y \otimes X \otimes Y \xrightarrow{\epsilon \otimes \text{id}_Y} Y$$

are identities. As mentioned in Remark 2.2.2, if Y exists, then it is unique so we will write $Y = X^\vee$ and call it the dual of X . Note that, as the conditions on X and $X^\vee$ are symmetric, we have $(X^\vee)^\vee = X$. In particular, we have both $X^\vee \otimes (-) \dashv X \otimes (-)$ and $X \otimes (-) \dashv X^\vee \otimes (-)$.

Now, assume further that the symmetric monoidal category $(\mathcal{C}, \otimes)$ is closed. That is, for any pair of objects $X, Y \in \mathcal{C}$, there is an internal hom object $\underline{\text{Hom}}(X, Y)$ such that for any object $Z \in \mathcal{C}$, there is an identification

$$\text{Hom}_{\mathcal{C}}(Z, \underline{\text{Hom}}(X, Y)) = \text{Hom}_{\mathcal{C}}(Z \otimes X, Y)$$

functorially on X, Y , and Z . (In fact, by the Yoneda lemma, the same identification holds when $\text{Hom}_{\mathcal{C}}$ itself is replaced by $\underline{\text{Hom}}$.) We have seen that Cat , $\widehat{\text{Cat}}^{\text{colim}}$, Pr^{L} and their stable variants are all closed symmetric monoidal categories. The observation to make here is that, when X admits a dual $X^\vee$, then both $X^\vee \otimes (-)$ and $\underline{\text{Hom}}(X, -)$ are right adjoints of $X \otimes (-)$ so they are the same. But then, evaluating at the tensor unit $1_{\mathcal{C}}$, we have

$$X^\vee = X^\vee \otimes 1_{\mathcal{C}} = \underline{\text{Hom}}(X, 1_{\mathcal{C}}).$$

More generally, one can deduce that $\underline{\text{Hom}}(X, Y) = X^\vee \otimes Y$.

We will prove later in Corollary 2.8.15 that when X is locally compact Hausdorff, the category $\text{Sh}(X; \text{Sp})$, as an object over spectra Sp , is self-dualizable. From this viewpoint, Lemma 2.5.11 is saying that, when $\mathcal{C} = \text{Sp}$,

$$\text{CSh}(X; \text{Sp}) = \text{Fun}^L(\text{Sh}(X), \text{Sp}) = \text{Fun}^L(\text{Sh}(X; \text{Sp}), \text{Sp})$$

is canonically dual to $\text{Sh}(X; \text{Sp})$ in $\text{Pr}_{st}^{\text{L}}$, or $\text{Sh}(X; \text{Sp})^\vee = \text{CSh}(X; \text{Sp})$. As a result, we have

Corollary 2.5.13. *Let X be a locally compact Hausdorff space and $\mathcal{C}$ be a stable category with small colimits. We have*

$$\text{CSh}(X; \mathcal{C}) = \text{CSh}(X; \text{Sp}) \otimes \mathcal{C}$$

where $\otimes$ is in general taken in $\widehat{\text{Cat}}^{\text{colim, st}}$ but is equal to being taken in $\text{Pr}_{st}^{\text{L}}$ if $\mathcal{C}$ is presentable.

Proof. By Lemma 2.5.11, $\text{CSh}(X; \mathcal{C}) = \text{Fun}^L(\text{Sh}(X); \mathcal{C})$. But the discussion before the lemma implies that

$$\text{Fun}^L(\text{Sh}(X); \mathcal{C}) = \text{Fun}^L(\text{Sh}(X; \text{Sp}); \mathcal{C}) = (\text{Sh}(X; \text{Sp}))^\vee \otimes \mathcal{C} = \text{CSh}(X; \text{Sp}) \otimes \mathcal{C}.$$

Note that we use the universal property of stabilization, Proposition 2.4.20, for the first equality. $\square$

Remark 2.5.14. We will see later that dualizability holds for $\text{Sh}(X; \mathcal{C})$ for a very general class of $\mathcal{C}$, by establishing the six-functor formalism, which in turn utilizes the Verdier duality that we will discuss shortly. In fact, the Verdier duality will also imply Corollary 2.5.13 for the case of sheaves. However, we mention this result here, since one can already establish the six-functor formalism for spectrum-valued sheaves directly from what we have already discussed regarding anima-valued sheaves and stabilization.

2.6. The Verdier duality. Consider a locally compact Hausdorff space X . On the level of abelian categories, there is a notion of the compactly supported sections of a sheaf F on an open set U which is given by

$$\Gamma_c(U; F) := \{s \in \Gamma(U; F) \mid \text{supp}(s) : \text{compact}\}.$$

If there is an inclusion $V \subseteq U$, then a section $s \in \Gamma_c(V; F)$ uniquely extends to a section $\tilde{s} \in \Gamma_c(U; F)$, since $\{V, U \setminus \text{supp}(s)\}$ forms a cover of U . That is, the assignment $U \mapsto \Gamma_c(U; F)$ forms a precosheaf. (Warning, this does not form a cosheaf on the abelian level. See Remark 2.6.1.)

We will define the higher categorical version of this assignment using what we have developed and show that, when X is a locally compact Hausdorff space and $\mathcal{C}$ a stable category which admits both limits and colimits, there is a (covariant) Verdier duality proved by Lurie [45, Section 5.5.5]

$$\mathrm{D}_{\mathcal{C}} : \mathrm{Sh}(X; \mathcal{C}) \xrightarrow{\sim} \mathrm{CSh}(X; \mathcal{C}),$$

which on objects is given by $F \mapsto (U \mapsto \Gamma_c(U; F))$. (We follow the treatment of Volpe here [66], who works some of the details claimed in [45].)

Remark 2.6.1. We note that the claim is not true at the abelian level and the statement only holds stably. For example, consider the case when $X = S^1$ and $F = C_{S^1}^{\ell c}$ given by locally constant functions. Then we have, in the abelian level, $\Gamma_c(S^1; C_{S^1}^{\ell c}) = \Gamma(S^1; C_{S^1}^{\ell c}) = \mathbb{R}$ but any open interval $I \subseteq S^1$ will have $\Gamma_c(I; C_{S^1}^{\ell c}) = 0$, so the cosheaf condition is not satisfied.

The Verdier duality will be defined in two steps. The first step is to show that the data of a sheaf on a locally compact Hausdorff space X can be recovered from its values on compact sets.

Definition 2.6.2 ([44, Section 7.3.4]). Let X be a locally compact Hausdorff space. We define $\mathcal{K}_X$ to be the poset of compact sets in X with respect to inclusion. We say that a functor $F : \mathcal{K}_X^{\mathrm{op}} \rightarrow \mathcal{C}$ is a $\mathcal{K}$ -sheaf on X if the following conditions are satisfied:

- (1) $F(\emptyset)$ is the final object.
- (2) For any $K_1, K_2 \in \mathcal{K}_X$, the commutative diagram

$$\begin{array}{ccc} \Gamma(K_1 \cup K_2; F) & \longrightarrow & \Gamma(K_1; F) \\ \downarrow & & \downarrow \\ \Gamma(K_2; F) & \longrightarrow & \Gamma(K_1 \cap K_2; F). \end{array}$$

is a pullback.

- (3) For any $K \in \mathcal{K}_X$, the canonical map

$$\mathrm{colim}_{K' \ni K} \Gamma(K'; F) \rightarrow \Gamma(K; F)$$

is invertible, where $K' \ni K$ runs over compact subsets K' such that $\mathrm{Int}(K') \supseteq K$.

We denote by $\mathrm{Sh}_{\mathcal{K}}(X; \mathcal{C})$ the category of $\mathcal{K}$ -sheaves.

Theorem 2.6.3 ([44, Theorem 7.3.4.9]). *Let X be a locally compact Hausdorff space and $\mathcal{C}$ a category admitting limits and colimits such that filtered colimits are left exact. Then, there is an equivalence $\mathrm{Sh}(X; \mathcal{C}) = \mathrm{Sh}_{\mathcal{K}}(X; \mathcal{C})$.¹⁵*

¹⁵See also [45, Lemma 5.5.5.3] regarding the presentability assumption in the cited reference.

Remark 2.6.4. Recall that one main property of a locally compact Hausdorff space is that, if K is a compact set and $W \supset K$ is open, there exists an open set U such that $\overline{U}$ is compact and $K \subseteq U \subseteq \overline{U} \subseteq W$.

We recall the proof for the reader's convenience: First, since X is a locally compact Hausdorff space, X is regular, meaning that any point p and closed set $C \not\ni p$ can be separated by open sets. A quick way to see this is by considering the one-point compactification $X \cup \{\infty\}$, which, as a compact Hausdorff space, is normal. But then

$$C \cup \{\infty\} = (X \cup \{\infty\}) \setminus (X \setminus C)$$

is a closed and thus compact in $X \cup \{\infty\}$, and thus can be separated from p .

The rest is then standard: Being locally compact and regular implies that, for any $x \in W$, there exists $U \ni x$ such that $\overline{U}$ is compact and contained in W . But then, K can be covered by finitely many such U and hence the result.

Proof. We construct a functor between the two categories by setting the poset $M_X := \text{Op}_X \cup \mathcal{K}_X$ within the power set of X . Let $i : \text{Op}_X \hookrightarrow M_X$ and $j : \mathcal{K}_X \hookrightarrow M_X$ be the inclusions. Precomposition and Kan extensions from Example 2.2.5 supply us two adjunctions

$$\begin{array}{ccccc} \text{PSh}(X; \mathcal{C}) & \xrightarrow{i_!} & \text{PSh}(M_X^{\text{op}}; \mathcal{C}) & \xrightarrow{j^*} & \text{PSh}(\mathcal{K}_X^{\text{op}}; \mathcal{C}) \\ & \perp & & \perp & \\ & \xleftarrow{i^*} & & \xleftarrow{j_*} & \end{array}$$

and we consider the functor $\theta := j^* i_! : \text{PSh}(X; \mathcal{C}) \rightarrow \text{PSh}(\mathcal{K}_X^{\text{op}}; \mathcal{C})$. Recall that j^* is just a restriction and, for a given $K \in \mathcal{K}_X$, what left Kan extension does is to approximate it by U and apply the functor to those approximations. More concretely, on objects, the functor θ is given by

$$\begin{aligned} \theta : \text{PSh}(X; \mathcal{C}) &\rightarrow \text{PSh}(\mathcal{K}_X; \mathcal{C}) \\ F &\mapsto \left(K \mapsto \text{colim}_{K \subseteq U} \Gamma(U; F) \right) \end{aligned}$$

The content of [44, Theorem 7.3.4.9] is that, after restricting to $\text{Sh}(X; \mathcal{C})$, θ takes value in $\text{Sh}_{\mathcal{K}}(X; \mathcal{C})$ and is an equivalence (with the restriction of $i^* j_*$ being the inverse). We will not explain all the details but simply check that a sheaf $\theta(F)$ is sent to a $\mathcal{K}$ -sheaf, and highlight how the conditions are used for this purpose.

For (1) of Definition 2.6.2, if $K = \emptyset$, then the diagram $\{U | K \subseteq U\}$ has $\emptyset$ as the initial object so $\theta(F)(\emptyset) = \Gamma(\emptyset; F) - *$ is terminal since F is a sheaf. For (2) of Definition 2.6.2, consider $K_\alpha \subseteq U_\alpha$ for $\alpha = 1, 2$ and recall that, since F is a sheaf, the diagram

$$\begin{array}{ccc} \Gamma(U_1 \cup U_2; F) & \longrightarrow & \Gamma(U_1; F) \\ \downarrow & & \downarrow \\ \Gamma(U_2; F) & \longrightarrow & \Gamma(U_1 \cap U_2; F). \end{array}$$

is a pullback. By the assumption that filtered colimits are left exact, the colimit of the diagram over all U_α containing K_α , $\alpha = 1, 2$, is a pullback and it gives the correct vertices

at K_1 and K_2 :

$$\begin{array}{ccc}
\operatorname{colim}_{U_\alpha \supseteq K_\alpha, \alpha=1,2} \Gamma(U_1 \cup U_2; F) & \xrightarrow{\quad} & \theta(F)(K_1) \\
\downarrow & & \downarrow \\
\theta(F)(K_2) & \xrightarrow{\quad} & \operatorname{colim}_{U_\alpha \supseteq K_\alpha, \alpha=1,2} \Gamma(U_1 \cap U_2; F).
\end{array}$$

That the other two are the desired objects follows from a cofinality argument. The case for the union is a direct consequence of Remark 2.6.4, as one can, for any $W \supset K_1 \cup K_2$, find U_α such that $K_\alpha \subseteq U_\alpha \subseteq \overline{U_\alpha} \subseteq W$. For the cofinality of the intersection, let $W \supseteq K_1 \cap K_2$; since $(K_1 \setminus W), (K_2 \setminus W)$ are two disjoint compact sets, by regularity, we can find relatively compact open sets V_α , $\alpha = 1, 2$ such that $V_\alpha \supseteq K_\alpha \setminus W$ and $V_1 \cap V_2 = \emptyset$. Set $U_\alpha = V_\alpha \cup W$, which contains K_α and satisfies $U_1 \cap U_2 = W$. (3) of Definition 2.6.2 also follows from a similar cofinality argument and the fact that sheaves satisfy their version of (3). $\square$

Definition 2.6.5. Let $F \in \operatorname{Sh}(X; \mathcal{C})$ and Z any closed subset of X . We define the sections of F supported on Z by

$$\Gamma_Z(X; F) := \operatorname{fib}(\Gamma(X; F) \rightarrow \Gamma(X \setminus Z; F)).$$

Note that, if $Z_1 \subseteq Z_2$, then we have a canonical map $\Gamma_{Z_1}(X; F) \rightarrow \Gamma_{Z_2}(X; F)$, i.e., the assignment is covariant. We can thus define the compactly supported sections of F over U , for $U \subseteq X$ open, as

$$\Gamma_c(U; F) := \operatorname{colim}_{K \subset U} \Gamma_K(X; F),$$

where the index $K \subset U$ runs over compact sets contained in U .

As the assignment $\Gamma_K(X; F)$ is covariant in K , there is also the functor

$$\begin{aligned}
(7) \quad & \operatorname{PSh}(X; \mathcal{C}) \rightarrow \operatorname{Fun}(\mathcal{K}_X; \mathcal{C}) \\
& F \mapsto (K \mapsto \Gamma_K(X; F)).
\end{aligned}$$

Now, observe that if $\mathcal{C}$ is a stable category with limits and colimits, then colimits are automatically left exact, as finite limits are also finite colimits. Furthermore, the opposite category $\mathcal{C}^{\operatorname{op}}$ will have the same properties, so we can apply Theorem 2.6.3 and obtain $\operatorname{CSh}(X; \mathcal{C}) = \operatorname{CSh}_{\mathcal{K}}(X; \mathcal{C})$, where $\operatorname{CSh}_{\mathcal{K}}(X; \mathcal{C}) := (\operatorname{Sh}_{\mathcal{K}}(X; \mathcal{C}^{\operatorname{op}}))^{\operatorname{op}}$. For this reason, the functor $\psi := i^* j_*$ will be given by

$$\begin{aligned}
(8) \quad & \psi : \operatorname{Fun}(\mathcal{K}_X; \mathcal{C}) \rightarrow \operatorname{Fun}(\operatorname{Op}_X, \mathcal{C}) \\
& G \mapsto \left(U \mapsto \operatorname{colim}_{K \subseteq U} \Gamma(K; G) \right).
\end{aligned}$$

Thus, we see that the composition of (7) and (8) induces a map

$$\begin{aligned}
& \operatorname{Fun}(\operatorname{Op}_X^{\operatorname{op}}; \mathcal{C}) \rightarrow \operatorname{Fun}(\operatorname{Op}_X, \mathcal{C}) \\
& F \mapsto (U \mapsto \Gamma_c(U; F)),
\end{aligned}$$

and we denote its restriction to $\operatorname{Sh}(X; \mathcal{C})$ by $\operatorname{D}_{\mathcal{C}}$ and call it the Verdier duality.

Theorem 2.6.6. *Let X be a locally compact Hausdorff space and $\mathcal{C}$ a stable category admitting both limits and colimits. The Verdier duality*

$$\begin{aligned} D_{\mathcal{C}} : \mathrm{Sh}(X; \mathcal{C}) &\xrightarrow{\sim} \mathrm{CSh}(X; \mathcal{C}) \\ F &\mapsto (U \mapsto \Gamma_c(U; F)) \end{aligned}$$

is an equivalence.

Proof. The details are explained in [66, Theorem 5.10] and we explain only a few steps as examples of how the argument goes. First, we show that if F is a sheaf, then $D_{\mathcal{C}}(F)$ is a cosheaf. By Theorem 2.6.3, it is enough to show that the assignment

$$K \mapsto \Gamma_K(X; F)$$

is a $\mathcal{K}$ -cosheaf. We check that the dual version of (2) of Definition 2.6.2 is satisfied and leave (1) and (3) to the reader: For K_{α} , $\alpha = 1, 2$, the claim is that the commutative diagram

$$\begin{array}{ccc} \Gamma_{K_1 \cap K_2}(X; F) & \longrightarrow & \Gamma_{K_1}(X; F) \\ \downarrow & & \downarrow \\ \Gamma_{K_2}(X; F) & \longrightarrow & \Gamma_{K_1 \cup K_2}(X; F). \end{array}$$

is a pushout. Since we are in the stable setting, it is the same as showing that it is a pullback. Recalling that $\Gamma_K(X; F) := \mathrm{fib}(\Gamma(X; F) \rightarrow \Gamma(X \setminus K; F))$, we can consider the corresponding pullback diagrams

$$\begin{array}{ccc} \Gamma(X; F) & \longrightarrow & \Gamma(X; F) \\ \downarrow & & \downarrow \\ \Gamma(X; F) & \longrightarrow & \Gamma(X; F) \end{array} \quad \begin{array}{ccc} \Gamma(X \setminus (K \cap K'); F) & \longrightarrow & \Gamma(X \setminus K_1; F) \\ \downarrow & & \downarrow \\ \Gamma(X \setminus K_2; F) & \longrightarrow & \Gamma(X \setminus (K_1 \cup K_2); F) \end{array}$$

and taking the fiber between them implies that the desired diagram is a pullback. The second step is to show that $D_{\mathcal{C}^{\mathrm{op}}}^{\mathrm{op}} : \mathrm{CSh}(X; \mathcal{C}) = (\mathrm{Sh}(X; \mathcal{C}^{\mathrm{op}}))^{\mathrm{op}} \rightarrow (\mathrm{CSh}(X; \mathcal{C}^{\mathrm{op}}))^{\mathrm{op}} = \mathrm{Sh}(X; \mathcal{C})$ is the inverse of $D_{\mathcal{C}}$. By symmetry, it is enough to show that $D_{\mathcal{C}^{\mathrm{op}}}^{\mathrm{op}} \circ D_{\mathcal{C}}$ is the identity.

We spell out what this amounts to. Let $G \in \mathrm{CSh}(X; \mathcal{C})$ be a cosheaf so, for an inclusion $V \subseteq U$, we have a morphism $G(V) \rightarrow G(U)$ in $\mathcal{C}$. To form the (co)Verdier duality, as in Definition 2.6.5, we use the notation

$$\begin{aligned} \Gamma^Z(X; G) &:= \mathrm{cof}(\Gamma(X \setminus Z; G) \rightarrow \Gamma(X; G)) \\ \Gamma^c(U; G) &:= \lim_{K \subseteq U} \Gamma^K(X; G) \end{aligned}$$

and so $D_{\mathcal{C}^{\mathrm{op}}}^{\mathrm{op}}(G)(U) = \Gamma^c(U; G)$. Thus,

$$(D_{\mathcal{C}^{\mathrm{op}}}^{\mathrm{op}} \circ D_{\mathcal{C}})(F)(U) = \Gamma^c(U; D_{\mathcal{C}}(F)) = \lim_{K \subseteq U} \Gamma^K(X; D_{\mathcal{C}}(F)),$$

and $\Gamma^K(X; D_{\mathcal{C}}(F)) = \mathrm{cof}(\Gamma(X \setminus K; D_{\mathcal{C}}(F)) \rightarrow \Gamma(X; D_{\mathcal{C}}(F)))$. The claim, which is the main computation of [66, Theorem 5.10], is that there is a fiber sequence

$$\Gamma_c(X \setminus K; F) \rightarrow \Gamma_c(X; F) \rightarrow \Gamma(K; F)$$

so $\Gamma^K(X; D_{\mathcal{C}}(F)) = \Gamma(K; F)$. We remark that while the computation is a complicated tracing of pullback diagrams in the cited proof, the answer is in fact expected, either from

the existence of the six-functor formalism or from the “exactness” property of compactly supported cohomology. In sum, we have

$$D_{\mathcal{C}^{\text{op}}}^{\text{op}}(D_{\mathcal{C}}(F))(U) = \lim_{K \subseteq U} \Gamma(K; F) = \lim_{K \subseteq U} \operatorname{colim}_{V \supseteq K} \Gamma(V; F).$$

The proof is then concluded by the next lemma. $\square$

Warning 2.6.7 ([47, Tag 06PV]). It is not true that $\mathcal{C}^{\text{op}}$ is always presentable when $\mathcal{C}$ is presentable. Thus, it is important that the construction of $D_{\mathcal{C}}$ does not depend on $\mathcal{C}$ being presentable but only that it has both limits and colimits.

Lemma 2.6.8. *Let X be a locally compact Hausdorff space. For a sheaf $F \in \operatorname{Sh}(X; \mathcal{C})$, the canonical map*

$$\Gamma(X; F) \rightarrow \lim_K \operatorname{colim}_{U \supseteq K} \Gamma(U; F),$$

where the iterated indices first run over compact sets K and then open sets U containing K , is an equivalence.

Proof. Set $\mathcal{U} := \{V | \bar{V} \text{ is compact}\}$ and we notice that, since X is locally compact, it forms a cover so

$$\Gamma(X; F) \xrightarrow{\sim} \lim_{V \in \mathcal{U}} \Gamma(V; F).$$

On the other hand, by Remark 2.6.4, compact sets of the form $\bar{V}$ are cofinal among compact sets K , and open sets U such that $\bar{U}$ is compact are cofinal among open sets $U \supseteq K$, so we have a sequence of maps

$$\Gamma(X; F) \rightarrow \lim_K \operatorname{colim}_{U \supseteq K} \Gamma(U; F) \xrightarrow{\sim} \lim_{V \in \mathcal{U}} \operatorname{colim}_{U \supseteq \bar{V}, U \in \mathcal{U}} \Gamma(U; F) \rightarrow \lim_{V \in \mathcal{U}} \Gamma(V; F)$$

whose composition is an equivalence. It then suffices to show that the last arrow

$$\lim_{V \in \mathcal{U}} \operatorname{colim}_{U \supseteq \bar{V}, U \in \mathcal{U}} \Gamma(U; F) \rightarrow \lim_{V \in \mathcal{U}} \Gamma(V; F)$$

has an inverse. But for any $V \in \mathcal{U}$, using Remark 2.6.4 again, we can fix $W \in \mathcal{U}$ such that $V \subseteq \bar{V} \subseteq W$, and we have

$$\Gamma(W; F) \rightarrow \operatorname{colim}_{U \supseteq \bar{V}, U \subseteq W} \Gamma(U; F) = \operatorname{colim}_{U \supseteq \bar{V}, U \in \mathcal{U}} \Gamma(U; F) \rightarrow \Gamma(V; F)$$

and it gives the desired inverse by taking limit over V . $\square$

Before leaving this section, we mention that (1) of Approximate Definition 2.0.1 in particular requires the sheaf $*$ -functoriality to admit extra adjoints when the map is either proper or an open embedding. By utilizing the fact that the “naive” functoriality of sheaves and cosheaves has opposite handiness, we will show that Theorem 2.6.6 gives a nice proof of such facts.

For that purpose, recall that we have an identification $\operatorname{CSh}(X; \mathcal{C}) = (\operatorname{Sh}(X; \mathcal{C}^{\text{op}}))^{\text{op}}$. One consequence of the ‘op’ outside is that the pullback functor for cosheaves should be a right adjoint of the pushforward functor. For reasons that will be clear later, we will denote the adjunction by $f_!^{\text{co}} \dashv f^{!, \text{co}}$. Similar to the sheaf situation, the naive pushforward functor on pre-cosheaves restricts to a functor between cosheaves

$$f_!^{\text{co}} : \operatorname{CSh}(X; \mathcal{C}) \rightarrow \operatorname{CSh}(Y; \mathcal{C})$$

$$F \mapsto (V \mapsto F(f^{-1}(V))).$$

To understand the pullback $f^{!,\text{co}}$, let $f : X \rightarrow Y$ be a map, $G : \text{Op}_Y \rightarrow \mathcal{C}$ be a precosheaf and $U \subseteq X$ be an open set. Then, the pre-cosheaf pullback, which we will denote by $f^{!,\text{co},\text{pre}}$, is given by

$$(f^{!,\text{co},\text{pre}}G)(U) = \lim_{W \supset f(U)} G(W).$$

Here, W runs over open sets containing $f(U)$, and we recall that since G is now a cosheaf, so if $f(U) \subseteq W_1 \subseteq W_2$, then the direction of the arrow is $G(W_1) \rightarrow G(W_2)$ and the correct approximation is a limit, which implies that $f^{!,\text{co},\text{pre}}$ should be the right adjoint.

Warning 2.6.9. While it is true that $f^{!,\text{co}}$ is obtained by composing $f^{!,\text{co},\text{pre}}$ with cosheafification, we actually do not yet know that we can cosheafify. That is, up to taking an ‘op’, we in fact have not proven that the sheafification functor $(-)^{\dagger} : \text{PSh}(X; \mathcal{C}) \rightarrow \text{Sh}(X; \mathcal{C})$ exists, except in the case when $\mathcal{C} = \text{Ani}$ or $\mathcal{C} = \text{Sp}$, and we thus do not yet have $*$ -pullback in general for $\text{Sh}(X; \mathcal{C})$. This subtlety will be dealt with in Section 2.9 again using the Verdier duality. But we point out that the next two lemmas are fine: f_* is always defined naively regardless of the coefficient $\mathcal{C}$ and, when $j : U \hookrightarrow X$ is an open embedding, the naive pullback $j^{*,\text{pre}}$ already sends sheaves to sheaves so sheafification is not needed. The same holds for cosheaves.

We assume again that the coefficient category $\mathcal{C}$ is a stable category with both small limits and colimits.

Lemma 2.6.10. *Let $f : X \rightarrow Y$ be a continuous map between locally compact Hausdorff spaces. Then, there is a natural transformation $f_!^{\text{co}} D_X \rightarrow D_Y f_*$. If $f = p$ is proper, then the natural transformation is an isomorphism.*

Proof. Let $F \in \text{Sh}(X; \mathcal{C})$ be a sheaf. Expanding the definition, we see that for $V \subseteq Y$ open,

$$((f_!^{\text{co}}(D_X F))(V) = (D_X(F))(f^{-1}(V)) = \Gamma_c(f^{-1}(V); F) = \text{colim}_{K \subseteq f^{-1}(V)} \Gamma_K(f^{-1}(V); F),$$

and

$$((D_Y(f_*F))(V) = \Gamma_c(V, (f_*F)) = \text{colim}_{L \subseteq V} \Gamma_L(V; f_*F) = \text{colim}_{L \subseteq V} \Gamma_{f^{-1}(L)}(f^{-1}(V); F).$$

Note that we use the fact that for any closed $Z \subseteq Y$, we have $\Gamma_Z(Y; f_*F) = \Gamma_{f^{-1}(Z)}(f^{-1}(Y); F)$ for the last equality on the second line. Indeed, by the definition, $\Gamma_Z(Y; f_*F)$ is the fiber of $\Gamma(Y; f_*F) \rightarrow \Gamma(Y \setminus Z; f_*F)$. But the morphism can be identified as $\Gamma(f^{-1}(Y); F) \rightarrow \Gamma(f^{-1}(Y) \setminus f^{-1}(Z); F)$, since for any sets $S_1, S_2 \subseteq Y$, we have $f^{-1}(S_1 \setminus S_2) = f^{-1}(S_1) \setminus f^{-1}(S_2)$.

Now, if $K \subseteq X$ is compact, then $f(K) \subseteq Y$ is compact and we have $K \subseteq f^{-1}(f(K))$. The canonical map $f_!^{\text{co}} D_X \rightarrow D_Y f_*$ is thus induced by

$$\Gamma_K(f^{-1}(V); F) \rightarrow \Gamma_{f^{-1}(f(K))}(f^{-1}(V); F).$$

Now, when $f = p$ is proper, any compact set L pulls back to a compact set $p^{-1}(L)$, which makes the second diagram a subdiagram of the first and this defines the inverse map $D_Y p_* \rightarrow p_!^{\text{co}} D_X$. The conclusion about the right adjoint of p_* follows from the observation before the lemma. $\square$

Lemma 2.6.11. *Let $j : U \hookrightarrow X$ be an open immersion of locally compact Hausdorff spaces. Then, there is a natural isomorphism $j^{!,\text{co}} D_X = D_U j^*$.*

Proof. Let $F \in \mathrm{Sh}(X; \mathcal{C})$ be a sheaf and $V \subseteq U$ be open. We again expand the definition and see

$$(j^{!,\mathrm{co}}(D_X F))(V) = (D_X(F))(V) = \Gamma_c(V; F),$$

and $(D_U(j^*F))(V) = \Gamma_c(V; j^*F)$. Tracing the definition, let $Z \subseteq V$ be closed and consider the fiber sequence

$$\Gamma_Z(V; j^*F) \rightarrow \Gamma(V; j^*F) \rightarrow \Gamma(V \setminus Z; j^*F).$$

But j^* only restricts the open sets which F takes value so $\Gamma_Z(V; j^*F) = \Gamma_Z(V; F)$ which implies $\Gamma_c(V; j^*F) = \Gamma_c(V; F)$. $\square$

Remark 2.6.12. We would like to conclude by Lemma 2.6.10 and Lemma 2.6.11, that when p is proper, the functor p_* admits a right adjoint $p^!$, which corresponds to $p^{!,\mathrm{co}}$ under the Verdier duality, and similarly, when j is open, j^* has a left adjoint $j_!$, which corresponds to $j_!^{\mathrm{co}}$ under Theorem 2.6.6. However, because of Warning 2.6.9, these extra adjunctions are not yet available to us without extra work. Nevertheless, we will assume their existence, and continue the discussion, as the author believes that it is pedagogically suitable.

Nevertheless, for readers who might worry about having a circular argument, we clarify the order of logic here: One first establishes the existence of $p^!$ and $j_!$ for the case of spectra-valued sheaves. We will explain how to obtain them in Section 2.9, more precisely in Corollary 2.9.2. With the existence of $p^!$ and $j_!$ for spectra-valued sheaves, we will be able to deduce, from the existence of the six-functor formalism, that $\mathrm{Sh}(X; \mathrm{Sp})$ is in fact self-dual in Corollary 2.8.15. Then, a trick of scalar extension, which again will be explained in Section 2.9, more precisely in Corollary 2.9.6, will imply the existence of f^* , $p^!$, and $j_!$ for any stable coefficient $\mathcal{C}$ with both small limits and colimits.

2.7. Category of correspondences.

Assumption 2.7.1. To keep the exposition streamlined, we assume throughout this section and the next (Section 2.8) the existence of the sheaf $*$ -functoriality, as well as the existence of $p^!$ and $j_!$. We remark that, in the case of spectra-valued sheaves, the $*$ -functoriality has already been established in Corollary 2.4.23, and $p^!$ and $j_!$ can be constructed independently of the discussion in the next two sections. Readers concerned about the logical order are referred to Warning 2.6.9 and Remark 2.6.12 for further details.

We remark that, we essentially have all the ingredients needed to produce the promised six-functor in Approximate Definition 2.0.1. But let us try to interpret what we have achieved: For a proper map $p : Y \rightarrow X$, the $*$ -pushforward p_* has an extra right adjoint $p^!$, or $p^* \dashv p_* \dashv p^!$. A better, more “internal” way to say it, as hinted at in Remark 2.2.23, is that $p^* \dashv p_*$ is an adjunction internal in $\mathbf{Pr}^{\mathbf{L}}$. Here, we use the bold font to emphasize that we now treat $\widehat{\mathbf{Cat}}$ as an $(\infty, 2)$ -category (of large $(\infty, 1)$ -categories) by allowing natural transformations that are not necessarily invertible, and we view $\mathbf{Pr}^{\mathbf{L}}$ as the non-full subcategory consisting of those objects and 1-morphisms in $\mathbf{Pr}^{\mathbf{L}}$.¹⁶

Notation 2.7.2. We will use LCH to denote the $(1, 1)$ -category of locally compact Hausdorff spaces with continuous maps, $\mathrm{LCH}_P \subseteq \mathrm{LCH}$ the wide subcategory¹⁷ consisting of only proper maps, and $\mathrm{LCH}_I \subseteq \mathrm{LCH}$ the wide subcategory consisting of open embeddings.

¹⁶An equivalent way to think of an $(\infty, 2)$ -category $\mathcal{C}$ other than having non-invertible 2-morphisms is that, for any two objects x, y , $\mathrm{Hom}_{\mathcal{C}}(x, y)$ should be an $(\infty, 1)$ -category.

¹⁷Meaning it contains all the objects.

To summarize, for a stable category $\mathcal{C}$ with small limits and colimits, what we have is thus a **2**-functor between **2**-categories

$$\begin{aligned} \mathrm{Sh}^*|_{\mathrm{LCH}_P} : \mathrm{LCH}_P^{\mathrm{op}} &\rightarrow \mathbf{Pr}^{\mathbf{L}} \\ X &\mapsto \mathrm{Sh}(X; \mathcal{C}) \\ (f : Y \rightarrow X) &\mapsto (f^* : \mathrm{Sh}(Y; \mathcal{C}) \rightarrow \mathrm{Sh}(X; \mathcal{C})), \end{aligned}$$

where we view LCH_P as a 2-category with only boring 2-morphisms, such that, for any $p : Y \rightarrow X$ in LCH_P , the morphism $\mathrm{Sh}^*(p) := p^*$ is internally a left adjoint in $\mathbf{Pr}^{\mathbf{L}}$. Furthermore, for a pullback diagram

$$\begin{array}{ccc} X' & \xrightarrow{q'} & X \\ \downarrow p' & & \downarrow p \\ Y' & \xrightarrow{q} & Y, \end{array}$$

the natural transformation we considered in Theorem 2.3.19

$$q^*p_* \rightarrow q^*p_*q'_*q'^* = q^*q_*p'_*q'^* \rightarrow p'_*q'^*$$

can be discussed also internally in $\mathbf{Pr}^{\mathbf{L}}$, and $\mathrm{Sh}^*|_{\mathrm{LCH}_P}$ makes it an isomorphism.¹⁸

As we have been emphasizing throughout this course, building a functor in the higher categorical setting is hard and one can imagine that building a 2-functor is surely even harder. As usual, the way to get around it is to somehow “enlarge” the ambient category in a universal way and reduce the question of building functors into verifying properties. In the situation at hand, what we have is a category $\mathcal{C}$ with finite limits, and what we would like to have is some kind of completion $\iota : \mathcal{C} \rightarrow \mathrm{Corr}(\mathcal{C})$, whose notation we will explain shortly, such that objects in $\mathrm{Corr}(\mathcal{C})$ should be the same as $\mathcal{C}$ but,

- (1) for any morphism $\alpha : x \rightarrow y$, we want to create for $\iota(\alpha)$ a right adjoint $\iota(\alpha)^R$, and
- (2) for any pullback diagram

$$\begin{array}{ccc} y' & \xrightarrow{\beta'} & y \\ \downarrow \alpha' & & \downarrow \alpha \\ x' & \xrightarrow{\beta} & x, \end{array}$$

as right adjoints exist, there is a canonical 2-morphism $\iota(\beta')\iota(\alpha')^R \rightarrow \iota(\alpha)\iota(\beta)$, and we ask it to be invertible.

The goal of this section is thus to first explain what $\mathrm{Corr}(\mathcal{C})$ looks like and describe its universal property. Secondly, one notices that if we restrict to the dual situation LCH_I where we only consider open embeddings, we will have a dual theory. However, the full six-functor formalism includes not just both situations but arbitrary continuous maps, though not much more in an appropriate sense. Such a complication forces us to consider an enhancement $\mathrm{Corr}(\mathcal{C})_{I,P}$, whose universal property we will state and comment on the significance of the subtlety. We now begin our discussion of $\mathrm{Corr}(\mathcal{C})$ with a definition. Our main reference here is [61].

¹⁸We will only use the bold font this one time to emphasize the subtlety, but otherwise abuse the notation and keep the simpler Cat . Similarly, we will stop keeping track of the categorical complexity and simply call a 2-functor a functor and trust that the reader can figure out the correct statement within the context.

Definition 2.7.3 ([61, Definition 3.4.1]). Let $\mathcal{D}$ be a 2-category. We say a commutative diagram

$$\begin{array}{ccc} d' & \xrightarrow{\alpha'} & d \\ \downarrow \beta' & & \downarrow \beta \\ e' & \xrightarrow{\alpha} & e. \end{array}$$

in $\mathcal{D}$ is vertically right adjointable if the following conditions hold:

- (1) The maps β and β' admit right adjoints β^R and β'^R .
- (2) The 2-morphism

$$\alpha' \beta'^R \rightarrow \beta^R \beta \alpha' \beta'^R = \beta^R \alpha \beta' \beta'^R \rightarrow \beta^R \alpha$$

built from the unit $\text{id}_d \rightarrow \beta^R \beta$ and the counit $\beta' \beta'^R \rightarrow \text{id}_{e'}$, is an isomorphism.

It is horizontally right adjointable if we instead reverse the roles of α , α' and β , β' in the definition. The diagram is right adjointable if it is both vertically and horizontally right adjointable.

Definition 2.7.4 ([61, Definition 3.4.5], Beck-Chevalley conditions). Let $\mathcal{C}$ be a category with pullbacks and $\mathcal{D}$ be a 2-category. We say that a functor $F : \mathcal{C} \rightarrow \mathcal{D}$ satisfies the left Beck-Chevalley condition if any pullback diagram

$$\begin{array}{ccc} y' & \xrightarrow{\beta'} & y \\ \downarrow \alpha' & & \downarrow \alpha \\ x' & \xrightarrow{\beta} & x \end{array}$$

in $\mathcal{C}$ is sent to a right adjointable diagram

$$\begin{array}{ccc} F(y') & \xrightarrow{F(\beta')} & F(y) \\ \downarrow F(\alpha') & & \downarrow F(\alpha) \\ F(x') & \xrightarrow{F(\beta)} & F(x) \end{array}$$

in $\mathcal{D}$.

Theorem 2.7.5 ([61, Theorem 3.4.18], [22, Ch.7, Theorem 3.2.2]). *Let $\mathcal{C}$ be a category with pullbacks. Then there is $\iota : \mathcal{C} \rightarrow \text{Corr}(\mathcal{C})$ exhibiting $\text{Corr}(\mathcal{C})$ as the universal 2-category that satisfies the left Beck-Chevalley condition. More precisely, for any 2-category $\mathcal{D}$, restriction along ι induces an inclusion*

$$\text{Hom}_{2\text{Cat}}(\text{Corr}(\mathcal{C}), \mathcal{D}) \hookrightarrow \text{Hom}_{2\text{Cat}}(\mathcal{C}, \mathcal{D})$$

with the image given by functors satisfying the left Beck-Chevalley condition.

We give an “approximate construction”¹⁹ of $\text{Corr}(\mathcal{C})$ and the map $\iota : \mathcal{C} \rightarrow \text{Corr}(\mathcal{C})$.

¹⁹For example, we ignore the issue on how to cut the size of $\text{Hom}_{\text{Corr}(\mathcal{C})}(-, -)$ so that $\text{Corr}(\mathcal{C})$ is locally small. See [61, Section 3.1] for the actual construction and, for this particular issue, see [61, Remark 3.1.10].

Approximate Definition 2.7.6. Let $\mathcal{C}$ be a category with pullbacks. The 2-category $\text{Corr}(\mathcal{C})$ consists of the same objects as $\mathcal{C}$. Let $x, y \in \mathcal{C}$. Morphisms from x to y are given by span diagrams of the form

$$\begin{array}{ccc} & u & \\ g \swarrow & & \searrow f \\ x & & y \end{array}$$

where $u \in \mathcal{C}$ is some object in $\mathcal{C}$ and f, g are morphisms in $\mathcal{C}$. Given two span diagrams, $x \xleftarrow{g} u \xrightarrow{f} y$ and $y \xleftarrow{h} v \xrightarrow{k} z$, their composition is given by the diagram

$$\begin{array}{ccccc} & & u \times_y v & & \\ & h' \swarrow & & \searrow f' & \\ & u & & v & \\ g \swarrow & & & & \searrow h \\ x & & y & & z \end{array}$$

where the middle diamond is the pullback diagram by f and h . That is,

$$(y \xleftarrow{h} v \xrightarrow{k} z) \circ (x \xleftarrow{g} u \xrightarrow{f} y) = x \xleftarrow{h' \circ g} u \times_y v \xrightarrow{k \circ f'} z.$$

Since $\mathcal{C}$ is a 1-category, this is all it should see and we can define the functor $\iota : \mathcal{C} \rightarrow \text{Corr}(\mathcal{C})$ by the identity on objects and for a morphism $f : x \rightarrow y$, its image $\iota(f)$ is given by

$$\begin{array}{ccc} & x & \\ \parallel \swarrow & & \searrow f \\ x & & y, \end{array}$$

or simply $\iota(f) = (x = x \xrightarrow{f} y)$.

Lastly, for two morphisms with the same domain and target, the 2-morphisms, from $x \xleftarrow{g'} u' \xrightarrow{f'} y$ to $x \xleftarrow{g} u \xrightarrow{f} y$, are given by commutative diagrams of the form

$$\begin{array}{ccc} & u' & \\ g' \swarrow & \downarrow & \searrow f' \\ & u & \\ g \swarrow & & \searrow f \\ x & & y. \end{array}$$

We now point out that, although we have only defined the notion of adjoint for functors in Definition 2.2.1, the same definition holds in any 2-category: The right adjoint of $f : x \rightarrow y$ is a morphism $f^R : y \rightarrow x$ of the reversed direction such that there exist a pair of 2-morphisms,

$$\epsilon : \text{id}_x \rightarrow f^R f \text{ and } \eta : f f^R \rightarrow \text{id}_y$$

such that the same triangle equality holds. For example, the composition

$$f = f \circ \text{id}_x \xrightarrow{\text{id}_f \circ \eta} f \circ (f^R \circ f) = (f \circ f^R) \circ f \xrightarrow{\epsilon \circ \text{id}_f} \text{id}_y \circ f = f$$

should be equal to id_f . Here, we use the notation $\circ_h$ to emphasize that they are “horizontal composition”. Note that in a (1-)category $\mathcal{C}$, since all 2-morphisms are invertible, asking for $f : x \rightarrow y$ to have a right adjoint is necessarily asking it to be invertible. So, this notion needs non-invertible 2-morphisms to be interesting. Furthermore, we want space for us to produce an arrow that “goes the wrong way”. The last piece of the definition gives us both.

Lemma 2.7.7. *Let $\mathcal{C}$ be a category. We have*

- (1) ([61, Proposition 3.3.1]) *For a morphism $f : x \rightarrow y$, $\iota(f) = (x = x \xrightarrow{f} y)$ has a right adjoint $\iota(f)^R = (y \xleftarrow{f} x = x)$, i.e., the diagram*

$$\begin{array}{ccc} & x & \\ f \swarrow & & \searrow \\ y & & x. \end{array}$$

- (2) ([61, Remark 3.1.14]) *The assignment $\iota(f)^R : \mathcal{C}^{\text{op}} \rightarrow \text{Corr}(\mathcal{C})$ organizes to a functor.*

Proof. We will only supply the unit and counit for the adjunction $\iota(f) \dashv \iota(f)^R$. First, one computes that $\iota(f)^R \circ \iota(f)$ is given by the diagram

$$\begin{array}{ccccc} & & x \times_y x & & \\ & p_1 \swarrow & & \searrow p_2 & \\ & x & & x & \\ & \parallel \swarrow & f \searrow & f \swarrow & \parallel \searrow \\ x & & y & & x, \end{array}$$

i.e., the self-fiber product and its natural projection $x \xleftarrow{p_1} x \times_y x \xrightarrow{p_2} x$. Similarly, we have $\iota(f) \circ \iota(f)^R = (x \xleftarrow{f} y \xrightarrow{f} x)$. The counit $\epsilon : \iota(f) \circ \iota(f)^R \rightarrow \text{id}_y$ is then given by the commutative diagram

$$\begin{array}{ccc} & x & \\ f \swarrow & & \searrow f \\ & y & \\ f \swarrow & & \searrow f \\ y & & y. \end{array}$$

and the unit $\eta : \text{id}_x \rightarrow \iota(f)^R \iota(f)$ by

$$\begin{array}{ccc} & x & \\ \parallel \swarrow & \Delta \downarrow & \parallel \searrow \\ & x \times_y x & \\ p_1 \swarrow & & \searrow p_2 \\ x & & x, \end{array}$$

where $\Delta : x \rightarrow x \times_y x$ is the diagonal map. □


Before discussing the general situation, we further assume that $\mathcal{C}$ has finite limits so there is a Cartesian monoidal structure $(\mathcal{C}, \times)$ given by taking products. The 2-category $\text{Corr}(\mathcal{C})$ thus can be equipped with a monoidal structure as well [61, Remark 3.2.5]. On objects, it is also given by $(X, Y) \mapsto X \times Y$. The amazing fact is that, with this monoidal structure, every object is self-dual, in the sense of Definition 2.5.12.

Lemma 2.7.8 ([61, Proposition 3.3.3]). *Let $x \in \mathcal{C}$ be an object. In $\text{Corr}(\mathcal{C})$, x is self-dual, i.e., $x^\vee = x$. More precisely, the unit and counit are given by*


$$\eta = (* \xleftarrow{a} x \xrightarrow{\Delta} x \times x) \text{ and } \epsilon = (x \times x \xleftarrow{\Delta} x \xrightarrow{a} *),$$

where $a : x \rightarrow *$ is the map to the final object and $\Delta : x \rightarrow x \times x$ is the diagonal.

Proof. We want to check, for example, $(\text{id}_x \times \epsilon) \circ (\eta \times \text{id}_x) = \text{id}_x$. Before composition, they are given by, in the order of $(\eta \times \text{id}_x)$ and then $(\text{id}_x \times \epsilon)$,




Thus, their composition is given by



but $\Delta \circ p_1 = \Delta \circ p_2 = \text{id}_x$ and thus the result. $\square$

Without going into too much detail, we remark on what extra enhancements are needed in order to work with the topology situation. We refer to [16] for interested readers. The thing we want to remark on is that, by considering open embeddings LCH_I , we want maps $\iota(f)$ to formally have left adjoints $\iota(f)^L$ instead of right adjoints. On the level of morphisms, we still use the trick of formally reversing the arrow $\iota(f)^L = (y \xleftarrow{f} x = x)$. While it might look confusing at first, we note that the reason why this works is that the notion of adjoint depends on what the 2-morphisms are, and to have them be the left adjoints, what has to be done is to reverse the central arrow



To have both adjointability, along with compatibility between them, i.e., (4) of Theorem 2.8.5, the construction is quite simple.

Approximate Definition 2.7.9 ([16, Construction 4.12]). Under certain compatibility conditions between I and P (Definition 2.8.3), the category $\text{Corr}(\mathcal{C})_{I,P}$ has the same objects and morphisms as $\text{Corr}(\mathcal{C})$ but the 2-morphisms, from $x \xleftarrow{g'} u' \xrightarrow{f'} y$ to $x \xleftarrow{g} u \xrightarrow{f} y$, are given by diagrams of the form

$$\begin{array}{ccccc}
 & & u' & & \\
 & \swarrow g' & \uparrow j & \searrow f' & \\
 x & & w & & y \\
 & \swarrow f & \downarrow p & \searrow g & \\
 & & u & &
 \end{array}$$

where $j \in I$ and $p \in P$.

Remark 2.7.10. This construction enjoys a generalized universal property, conceptually similar yet more complicated to state than the one described in Theorem 2.7.5. This is one of the main theorems [16, Theorem A] of the paper. Similarly, as explained in [16, Theorem B], a nice symmetric monoidal structure can be put on $\text{Corr}_{I,P}(\mathcal{C})$.

2.8. Six-functor formalism in topology: Spectra coefficient. Assuming the definition of the category of correspondences $\text{Corr}(\mathcal{C})$, we can now give the official definition of a six-functor formalism.

Definition 2.8.1 ([30, Definition 3.1.1, 3.2.1]). Let $\mathcal{C}$ be a category with pullbacks and view $\text{Corr}(\mathcal{C})$ as a symmetric monoidal 1-category²⁰ with the monoidal structure described before Lemma 2.7.8. A three-functor formalism is a symmetric monoidal functor $D : \text{Corr}(\mathcal{C}) \rightarrow (\widehat{\text{Cat}}, \times)$. A six-functor formalism is a factorization of a three-functor formalism to

$$D : \text{Corr}(\mathcal{C}) \rightarrow (\text{Pr}^{\text{L}}, \otimes).$$

Remark 2.8.2. There are a few points we want to mention regarding the definition of six-functor formalism.

- (1) While we define $\text{Corr}(\mathcal{C})$ as a 2-category in Approximate Definition 2.7.6, in practice, we care only about its underlying 1-category (obtained by forgetting all non-invertible 2-morphisms). The logic is the following: We check the universal property (concretely as Proposition 2.8.5), which only holds true as a 2-category, so we can extend a given functor $D^* : \mathcal{C}^{\text{op}} \rightarrow (\text{Pr}^{\text{L}}, \otimes)$ to $D : \text{Corr}(\mathcal{C}) \rightarrow (\text{Pr}^{\text{L}}, \otimes)$. But then, restrict D to the underlying 1-category and forget about the 2-morphisms.
- (2) Instead of considering all morphisms in $\mathcal{C}$, it is often more natural to pick a family E of “!-able morphisms”, and consider the pair $(\mathcal{C}, E)$ as the “geometric setup”. We will not go into details of it, as we see in our case all maps can be made !-able, but we will include it in the statements for completeness. Readers interested only in the topological situation can simply take E as all 1-morphisms in $\mathcal{C}$.
- (3) The definition given in [30, Definition 3.2.1] is in fact slightly more general, as they do not assume the categories to be presentable, but only the existence of right adjoints. Though, this is usually the case so we choose this presentation for simplicity.

²⁰See Remark 2.8.2 below.

Definition 2.8.3 ([30, Definition 3.3.2]). Let $\mathcal{C}$ be a category with pullbacks and E a family of !-able maps. A suitable decomposition consists of two families I, P of morphisms in E such that

- (1) I and P contain the isomorphisms and are stable under composition and pullbacks. Here, stable under pullbacks means that for any pullback diagram

$$\begin{array}{ccc} X' & \xrightarrow{g'} & X \\ \downarrow f' & & \downarrow f \\ Y' & \xrightarrow{g} & Y \end{array}$$

such that $f \in P$ (resp. $f \in I$), the map $f' \in P$ (resp. $f' \in I$).

- (2) Any map f in E can be decomposed as $f = p \circ j$ for some $p \in P$ and $j \in I$.
- (3) I and P are right cancellative. Given $X \xrightarrow{f} Y \xrightarrow{g} Z$, if g and $g \circ f$ are in P (resp. in I), then f is in P (resp. is in I).
- (4) Every morphism $f \in I \cap P$ is n -truncated for some $n \geq -2$ (possibly depending on f). Note that this is always satisfied if $\mathcal{C}$ is an ordinary category.

Lemma 2.8.4. *In the category LCH of locally compact Hausdorff spaces with morphisms given by continuous maps, the two families P of proper maps and I of open immersions form a suitable decomposition for LCH.*

Proof. Conditions (1) and (3) are straightforward. For (2), let $f : X \rightarrow Y$ be a map. We denote by $\mathbf{CHaus}$ the category of compact Hausdorff spaces; the inclusion $\mathbf{CHaus} \subseteq \mathbf{Top}$ admits a left adjoint $\beta : \mathbf{Top} \rightarrow \mathbf{CHaus}$, given explicitly by the Stone-Ćech compactification, and it has the property that the unit $\eta_X : X \rightarrow \beta X$ is an open immersion if and only if X is locally compact Hausdorff. Thus, we set $\overline{X}$ to be the pullback of $Y \hookrightarrow \beta Y$ along $\beta(f) : \beta X \rightarrow \beta Y$ and consider the following commutative diagram, decomposing $f = p \circ j$:

$$\begin{array}{ccccc} X & & & & Y \\ & \searrow j & & \nearrow p & \\ & \overline{X} & & & \\ \eta_X \swarrow & \downarrow \eta'_Y & & \downarrow \eta_Y & \\ & \beta X & \xrightarrow{\beta(f)} & \beta Y & \end{array}$$

Since Y is locally compact Hausdorff, η_Y is an open inclusion. Using (1) and then (3) for I , we conclude that $j : X \hookrightarrow \overline{X}$ is an open immersion. Since maps between compact Hausdorff spaces are proper, using (1) for P , we conclude that p is proper. $\square$

Proposition 2.8.5 ([30, Proposition 3.3.3]). *Let $(\mathcal{C}, E)$ be a geometric setup such that $\mathcal{C}$ admits pullbacks, let $I, P \subseteq E$ be a suitable decomposition of E and let $D^* : \mathcal{C}^{\text{op}} \rightarrow (\mathbf{Cat}, \times)$ be a symmetric monoidal functor. Assume that the following conditions are satisfied:*

- (1) *For every map $j : U \rightarrow X$ in I , the functor j^* admits a left adjoint $j_! : D(U) \rightarrow D(X)$ and it satisfies base-change and the projection formula.*
- (2) *For every map $p : X \rightarrow Y$ in P , the functor p^* admits a right adjoint $p_* : D(X) \rightarrow D(Y)$ and it satisfies base-change and the projection formula.*

(3) For any pullback diagram

$$\begin{array}{ccc} U' & \xrightarrow{j'} & X' \\ \downarrow p' & & \downarrow p \\ U & \xrightarrow{j} & X \end{array}$$

in $\mathcal{C}$ such that $j \in I$ and $p \in P$, the natural map $j_! f'_* \xrightarrow{\sim} f_* j'_!$ is an isomorphism of functors $D(U') \rightarrow D(X)$.

Then, D extends to a 3-functor formalism $D : \text{Corr}(\mathcal{C}, E)^\otimes \rightarrow \text{Cat}$ such that for every $j \in I$ the functor $j_!$ is left adjoint to j^* and for every $p \in P$ the functor $p_!$ is right adjoint to p^* .

We have proven the existence of $j_! \vdash j^*$ and $p^! \dashv p_*$ when j is an open immersion and p is proper in Lemma 2.6.10 and 2.6.11, for a very general class of coefficients. In fact, most of the conditions can be checked directly, possibly with the help of the Verdier duality from Theorem 2.6.6. However, to conclude proper base-change, we need to first deduce the basic case $\mathcal{C} = \text{Sp}$ from what we have recalled about anima-valued sheaves, before upgrading the coefficients by using a sheaf version of Corollary 2.5.13. We also notice that the discussion about Sh^* being a symmetric monoidal functor and the projection formula are easier in this case, since we have the Künneth formula for anima, Proposition 2.5.8.

We begin with conditions implied by the Verdier duality. For these conditions, we can work with a general coefficient $\mathcal{C}$ directly.

Lemma 2.8.6. *Let $\mathcal{C}$ be a stable category with small limits and colimits. Then, the functor*

$$\begin{aligned} \text{LCH} &\rightarrow \text{Pr}^L \\ X &\mapsto \text{Sh}(X; \mathcal{C}) \end{aligned}$$

satisfies base-change for open immersions. That is, if $j : U \hookrightarrow X$ is an open immersion and

$$\begin{array}{ccc} U' & \xrightarrow{j'} & X' \\ \downarrow f' & & \downarrow f \\ U & \xrightarrow{j} & X \end{array}$$

a pullback diagram, then the canonical map $f^ j_! \rightarrow j'_! f'^*$ is an isomorphism.*

Proof. Recall that there is an anti-equivalence of $\text{Pr}^L = (\text{Pr}^R)^{\text{op}}$ given by passing to the right adjoints. Since the notion of being isomorphic is preserved under equivalence, it is sufficient to check that the right adjoints $j^* f_* \leftarrow f'_* j'^* : \text{Sh}(U; \mathcal{C}) \rightarrow \text{Sh}(X'; \mathcal{C})$ are isomorphisms. But we recall that, in the topological case, $U' = f^{-1}(U)$ is simply the inverse image of U under f . Let $V \subseteq U$ and $F \in \text{Sh}(X; \mathcal{C})$; we can compute that

$$\Gamma(V; f'_* j'^* F) = \Gamma(f^{-1}(V); j'^* F) = \Gamma(f^{-1}(V); F) = \Gamma(V; f_* F) = \Gamma(V; j^* f_* F).$$

□

Lemma 2.8.7. *Condition (3) in Proposition 2.8.5 holds for the functor*

$$\begin{aligned} \text{LCH} &\rightarrow \text{Pr}^L \\ X &\mapsto \text{Sh}(X; \mathcal{C}) \end{aligned}$$

Proof. Let

$$\begin{array}{ccc} U' & \xrightarrow{j'} & X' \\ \downarrow p' & & \downarrow p \\ U & \xrightarrow{j} & X \end{array}$$

be a pullback diagram where j is an open immersion and p is proper, and we need to check that $j_! p'_* \rightarrow p_* j'_!$ is an isomorphism. By Lemma 2.6.10 and Lemma 2.6.11, under the Verdier duality in Theorem 2.6.6, it corresponds to $j'_!{}^{\text{co}} p'_!{}^{\text{co}} \rightarrow p_!{}^{\text{co}} j_!{}^{\text{co}}$ and it is an isomorphism by the functoriality of $!$ -pushforward for cosheaves. $\square$

Lemma 2.8.8. *Over the spectra, the functor*

$$\begin{aligned} \text{LCH} &\rightarrow \text{Pr}^{\text{L}} \\ X &\mapsto \text{Sh}(X; \text{Sp}) \end{aligned}$$

satisfies proper base-change for proper maps. That is, if $p : X \rightarrow Y$ is proper and

$$\begin{array}{ccc} X' & \xrightarrow{q'} & X \\ \downarrow p' & & \downarrow p \\ Y' & \xrightarrow{q} & Y \end{array}$$

a pullback diagram, then the canonical map $q^ p_* \rightarrow p'_* q'^*$ is an isomorphism.*

Proof. Over anima, Theorem 2.3.19 states that the canonical map $q^* p_* \xrightarrow{\sim} p'_* q'^*$ is an isomorphism. Similar to the case that a functor sends an invertible morphism to an invertible morphism, a 2-functor sends an invertible 2-morphism to an invertible 2-morphism. Thus, as we have the identification $\text{Sp}(\text{Sh}(X)) = \text{Sh}(X; \text{Sp})$ from Example 2.4.22, we only need to convince ourselves that the functor $\text{Sp}(-) : \text{Pr}^{\text{L}} \rightarrow \text{Pr}_{st}^{\text{L}}$ in fact enhances to a 2-functor. This can be done by recalling that $\text{Sp}(\mathcal{C})$ is constructed by a limit diagram natural in $\mathcal{C}$ or, more concisely, by using the identification $\text{Sp}(-) = \text{Sp} \otimes (-)$ from (4) of Proposition 2.5.5. $\square$

Now, we turn to the discussion of symmetric monoidal structure and projection formula. First, recall that Proposition 2.5.8 implies that over Ani , $\text{Sh}^* : \text{LCH} \rightarrow \text{Pr}^{\text{L}}$ is symmetric monoidal, and we notice that it naturally induces a symmetric monoidal structure on $\text{Sh}(X)$.

Definition 2.8.9. Denote by $\Delta : X \hookrightarrow X \times X$ the diagonal. We define a symmetric monoidal structure $\times$ on $\text{Sh}(X)$ by the composition

$$\text{Sh}(X) \times \text{Sh}(X) \rightarrow \text{Sh}(X \times X) \xrightarrow{\Delta^*} \text{Sh}(X).$$

With the notation of (4), for $F, G \in \text{Sh}(X)$, $F \times G := \Delta^*(F \boxtimes G)$. By the universal property of the Lurie tensor product, Proposition 2.5.2, this symmetric monoidal structure is equivalent to the map

$$(9) \quad \text{Sh}(X) \otimes \text{Sh}(X) = \text{Sh}(X \times X) \xrightarrow{\Delta^*} \text{Sh}(X).$$

Lemma 2.8.10. *Let $f : X \rightarrow Y$ be a map. The functor $f^* : \text{Sh}(Y) \rightarrow \text{Sh}(X)$ is symmetric monoidal.*

Proof. This follows from Proposition 2.5.8 applied to $(f \times f) \circ \Delta_X = \Delta_Y \circ f$. $\square$

Lemma 2.8.11. *Over the spectra, the functor*

$$\begin{aligned} \text{LCH} &\rightarrow \text{Pr}^{\text{L}} \\ X &\mapsto \text{Sh}(X; \text{Sp}) \end{aligned}$$

is symmetric monoidal. That is, for $X, Y \in \text{LCH}$, there is a canonical equivalence

$$\text{Sh}(X; \text{Sp}) \otimes \text{Sh}(Y; \text{Sp}) \xrightarrow{\sim} \text{Sh}(X \times Y; \text{Sp})$$

Furthermore, if $f : X' \rightarrow X$ and $g : Y' \rightarrow Y$ are two maps, we have $f^ \otimes g^* = (f \times g)^*$.*

Proof. The canonical map can be defined similarly as in the case of anima-valued sheaves. One can check directly that it is the same as stabilizing the construction there. Equivalently, it can be obtained from the anima case by tensoring with Sp , as (4) of Proposition 2.5.5 and Example 2.4.22 imply that

$$\text{Sh}(X; \text{Sp}) = \text{Sp}(\text{Sh}(X)) = \text{Sh}(X) \otimes \text{Sp}.$$

The latter two viewpoints imply the equivalence. □

One can likewise define a symmetric monoidal structure $\otimes$ on $\text{Sh}(X; \text{Sp})$ by

$$(10) \quad \text{Sh}(X; \text{Sp}) \otimes \text{Sh}(X; \text{Sp}) = \text{Sh}(X \times X; \text{Sp}) \xrightarrow{\Delta^*} \text{Sh}(X; \text{Sp}).$$

The anima version then implies its compatibility with $*$ -pullback.

Proposition 2.8.12. *Let $f : X \rightarrow Y$ be a map. Then, the functor*

$$f^* : \text{Sh}(Y; \text{Sp}) \rightarrow \text{Sh}(X; \text{Sp})$$

is symmetric monoidal.

Lemma 2.8.13. *Let $f : X \rightarrow Y$ be a map, $F \in \text{Sh}(X; \text{Sp})$, and $G \in \text{Sh}(Y; \text{Sp})$. We have the following projection formulas:*

(1) *Assume $f = j$ is an open embedding. Then, there is a canonical map*

$$j_!(F \otimes j^*G) \rightarrow (j_!F) \otimes G$$

which is an isomorphism.

(2) *Assume $f = p$ is proper. Then, there is a canonical map*

$$(p_*F) \otimes G \rightarrow p_*(F \otimes p^*G)$$

which is an isomorphism.

Proof. We show that the projection formulas are in fact a consequence of base change. We will explain the case of a proper map p . Consider the pullback diagram

$$\begin{array}{ccc} X & \xrightarrow{\Gamma_p} & X \times Y \\ \downarrow p & & \downarrow (p \times \text{id}) \\ Y & \xrightarrow{\Delta} & Y \times Y \end{array}$$

where $\Gamma_p : X \hookrightarrow X \times Y$ is the graph $\Gamma_p(x) := (x, p(x))$. We can thus compute that

$$\begin{aligned} (p_*F) \otimes G &= \Delta^*(p_*F \boxtimes G) = \Delta^*(p_* \otimes \text{id})(F \boxtimes G) = \Delta^*(p \times \text{id})_*(F \boxtimes G) \\ &\xrightarrow{\sim} p_*\Gamma_p^*(F \boxtimes G) = p_*(F \otimes p^*G). \end{aligned}$$

Here, the isomorphism $\xrightarrow{\sim}$ at the beginning of the second row is given by base change. Note also that we implicitly use Lemma 2.6.10, 2.8.11, and Corollary 2.9.2 here: In order to make sense of $p_* \otimes \text{id}$, p_* has to be a morphism in Pr^{L} . But we know in fact that $p_* \vdash p^*$ internally in Pr^{L} and it further implies that $p_* \otimes \text{id} \vdash p^* \otimes \text{id}$. As right adjoints are unique, $p^* \otimes \text{id} = p^* \otimes \text{id}^* = (p \times \text{id})^*$ then implies that $p_* \otimes \text{id} = (p \times \text{id})_*$. $\square$

Having checked all the conditions in Proposition 2.8.5, we obtain the six-functor formalism over spectra-valued sheaves.

Proposition 2.8.14. *The functor*

$$\begin{aligned} \text{Sh}^* : \text{LCH}^{\text{op}} &\rightarrow \text{Pr}^{\text{L}} \\ X &\mapsto \text{Sh}(X; \text{Sp}) \end{aligned}$$

is symmetric monoidal. The families P and I given by proper maps and open immersions satisfy the conditions in Proposition 2.8.5. As a result, it extends to a six-functor formalism $\text{Sh} : \text{Corr}(\text{LCH})^{\otimes} \rightarrow \text{Pr}^{\text{L}}$.²¹

Corollary 2.8.15. *Let $X \in \text{LCH}$. The category $\text{Sh}(X; \text{Sp}) \in \text{Pr}_{st}^{\text{L}}$ is self-dualizable.*

Remark 2.8.16. With the six-functor formalism for spectra-valued sheaves established, we notice that if $f : X \rightarrow Y$ is in LCH , then $D_Y f_! = f_!^{\text{co}} D_X$ and hence the choice of notation for the cosheaf functoriality. Indeed, $f_!$ is essentially constructed by systematically setting it as $f_! = p_* j_!$ where $p \in P$ and $j \in I$ are some factorization. But then, the two cases are identified in Lemma 2.6.10 and 2.6.11.

Proof. By Lemma 2.7.8, any object $X \in \text{LCH}$ is self-dual in $\text{Corr}(\text{LCH})^{\otimes}$. But dualizable objects are sent to dualizable objects, so $\text{Sh}(X; \text{Sp})$ being self-dual is implied directly by the existence of the six-functor formalism. $\square$

Remark 2.8.17. Corollary 2.8.15 in fact says something very concretely: Under the identification $\text{Sh}(X; \text{Sp}) \otimes \text{Sh}(X; \text{Sp}) = \text{Sh}(X \times X; \text{Sp})$, the unit η and counit ϵ witnessing the self-duality of $\text{Sh}(X; \text{Sp})$ are given by the compositions

$$\eta : \text{Sh}(*) \xrightarrow{a^*} \text{Sh}(X; \text{Sp}) \xrightarrow{\Delta_!} \text{Sh}(X \times X; \text{Sp})$$

and

$$\epsilon : \text{Sh}(X \times X; \text{Sp}) \xrightarrow{\Delta^*} \text{Sh}(X; \text{Sp}) \xrightarrow{a_!} \text{Sh}(*) .$$

Here, we use $a : X \rightarrow *$ to denote the map to a point. Furthermore, the triangle equalities in this case can be obtained by base-changing along the pullback diagram

$$\begin{array}{ccc} X & \xrightarrow{\Delta} & X \times X \\ \downarrow \Delta & & \downarrow \text{id} \times \Delta \\ X \times X & \xrightarrow{\Delta \times \text{id}} & X \times X \times X. \end{array}$$

Proof. The first statement follows from Corollary 2.9.6 above. For the second, it boils down to show that $f_{\text{Sp}}^* \otimes \text{id}_{\mathcal{C}} = f_{\mathcal{C}}^*$. As in the proof of Proposition 2.9.12, it's sufficient to check on $\text{Op}_X \times \mathcal{C}$, i.e., the maps have to match on objects of the form $h_U \otimes M$ for $M \in \mathcal{C}$. Since $f_{\mathcal{C}}^* = (-)_{\mathcal{C}}^{\dagger} \circ f_{\mathcal{C}}^{*, \text{pre}}$, one can check them separately. The claim on the presheaf level

²¹Implicitly, we are mapping to $\widehat{\text{Cat}}$ so we can discuss whether it factorizes to Pr^{L} . But, changing the universe does not affect the statement so Proposition 2.8.5 can be applied.

is straightforward, and the only non-trivial statement is that $(-)^{\dagger}_{\mathrm{Sp}} \otimes \mathrm{id}_{\mathcal{C}} = (-)^{\dagger}_{\mathcal{C}}$, which is proved in details in [66, Theorem 5.23]. We only recall that, the equivalence $\mathrm{Sh}(X; \mathcal{C}) = \mathrm{Sh}(X) \otimes \mathcal{C}$ goes to $\mathrm{Fun}^L(\mathrm{Sh}(X), \mathcal{C}) = \mathrm{CSh}(X; \mathcal{C}) = \mathrm{Sh}(X; \mathcal{C})$ so we would like to trace where the objects go.

For that, the goal is to show that $(h_U)^{\dagger}_{\mathrm{Sp}} \otimes M$ is the sheafification of $h_U \otimes M$. One thus recall that, in Lemma 2.5.11, the identification is given by

$$\begin{aligned} \mathrm{Fun}^L(\mathrm{Sh}(X), \mathcal{C}) &\rightarrow \mathrm{CSh}(X; \mathcal{C}) \\ F &\mapsto (U \mapsto F(h_U)) \end{aligned}$$

and the Verdier duality is given by

$$\begin{aligned} D_{\mathcal{C}} : \mathrm{Sh}(X; \mathcal{C}) &\rightarrow \mathrm{CSh}(X; \mathcal{C}) \\ F &\mapsto (U \mapsto \Gamma_c(U; F)). \end{aligned}$$

Composing the two, we see that the representative $h_U \otimes M \in \mathrm{Sh}(X; \mathcal{C})$ is sent to

$$(F \mapsto \Gamma_c(U; F) \otimes M),$$

and, once tracing through the chain

$$\mathrm{Sh}(X; \mathcal{C}) = \mathrm{CSh}(X; \mathcal{C}) = \mathrm{Fun}^L(\mathrm{Sh}(X), \mathcal{C}) = \mathrm{Fun}^L(\mathrm{Sh}(X), \mathrm{Sp}) \otimes \mathcal{C} = \mathrm{CSh}(X) \otimes \mathcal{C} = \mathrm{Sh}(X) \otimes \mathcal{C},$$

one will see that it matches the desired object. $\square$

2.9. The *-functoriality over general coefficients. We begin this section by verifying Assumption 2.7.1 for spectra-valued sheaves, thereby establishing the six-functor formalism over the sphere spectrum. Since the geometric objects of interest are often naturally defined over other coefficients. For example, in the case of D -modules, one must work over the complex numbers.²² We will verify the same assumption for all coefficients $\mathcal{C}$ for which Verdier duality holds. To begin, we must first establish the *-functoriality over a general coefficient category $\mathcal{C}$. For this purpose, we employ an extension-of-scalars argument, using the fact that $\mathrm{Sh}(X; \mathrm{Sp})$ is dualizable in Pr_{st}^L (Corollary 2.9.6).

Lemma 2.9.1. *Let $f : X \rightarrow Y$ be a continuous map and let $\mathcal{C}$ be a category with colimits. The cosheaf pushforward*

$$f_!^{\mathrm{co}} : \mathrm{CSh}(X; \mathcal{C}) \rightarrow \mathrm{CSh}(Y; \mathcal{C})$$

corresponds to

$$(-) \circ f^* : \mathrm{Fun}^L(\mathrm{Sh}(X); \mathcal{C}) \rightarrow \mathrm{Fun}^L(\mathrm{Sh}(Y); \mathcal{C})$$

under the identification $\mathrm{Fun}^L(\mathrm{Sh}(X); \mathcal{C}) = \mathrm{CSh}(X; \mathcal{C})$ of Lemma 2.5.11.

Proof. Recall that the identification $\mathrm{Fun}^L(\mathrm{Sh}(X); \mathcal{C}) \simeq \mathrm{CSh}(X; \mathcal{C})$ is given by sending a colimit-preserving functor $F : \mathrm{Sh}(X) \rightarrow \mathcal{C}$ to the cosheaf $\bar{F}(U) := F(h_U)$. One then computes that, for $V \in \mathrm{Op}_Y$,

$$(f_!^{\mathrm{co}} \bar{F})(V) = \bar{F}(f^{-1}V) = F(h_{f^{-1}V}),$$

and similarly

$$(\overline{F \circ f^*})(V) = F(f^* h_V) = F(h_{f^{-1}V}).$$

$\square$

²²For traditional coefficients such as this, or more generally ω -presentable coefficients, the approach presented here is not necessary. However, there are coefficients that arise naturally from geometry and do not fall into this class, e.g., the wild Betti coefficient [58], which appears naturally when considering symplectic cohomology [42].

Corollary 2.9.2. *Let $f : X \rightarrow Y$ be a continuous map and let $\mathcal{C}$ be presentable. Then the functor $f_!^{\text{co}} : \text{CSh}(X; \mathcal{C}) \rightarrow \text{CSh}(Y; \mathcal{C})$ admits a right adjoint*

$$f^{!, \text{co}} : \text{CSh}(Y; \mathcal{C}) \rightarrow \text{CSh}(X; \mathcal{C}).$$

In particular, if we further assume that $X \in \text{LCH}$ and $\mathcal{C}$ is a stable presentable category, then $\text{Sh}(X; \mathcal{C})$ is also presentable and we have

(1) *if $p : X \rightarrow Y$ is proper, then p_* admits a right adjoint*

$$p^! : \text{Sh}(Y; \mathcal{C}) \rightarrow \text{Sh}(X; \mathcal{C});$$

(2) *if $j : U \hookrightarrow X$ is an open embedding, then j^* admits a left adjoint*

$$j_! : \text{Sh}(U; \mathcal{C}) \rightarrow \text{Sh}(X; \mathcal{C}).$$

Proof. By part (3) of Proposition 2.5.5, $\text{Fun}^L(-, \mathcal{C})$ is an internal Hom in Pr^L . In particular, precomposing with f^* , the functor $(-) \circ f^*$ is a morphism in Pr^L and thus admits a right adjoint at the level of underlying categories. The corresponding functor $f_!^{\text{co}}$ therefore also admits a right adjoint. The conclusions about $j_!$ and $p^!$ then follow directly from Lemmas 2.6.10 and 2.6.11. $\square$

Remark 2.9.3. The same argument, using Proposition 2.5.5 and Lemma 2.5.11, also shows that when $\mathcal{C}$ is presentable (but not necessarily stable), both $\text{Sh}(X; \mathcal{C})$ and $\text{CSh}(X; \mathcal{C})$ are presentable.

With Corollary 2.9.2, the discussion in Sections 2.7 and 2.8 goes through, and we in particular obtain the fact that $\text{Sh}(X; \text{Sp})$ is self-dual in Pr_{st}^L (Corollary 2.8.15). We will use this to establish the existence of the $*$ -functoriality outside of the presentable setting.

Lemma 2.9.4. *Let X be a locally compact Hausdorff space and let $\mathcal{C}$ be a stable category with both limits and colimits. Then*

$$\text{Sh}(X; \mathcal{C}) = \text{Sh}(X; \text{Sp}) \otimes \mathcal{C}.$$

Proof. Since $\mathcal{C}$ is stable and admits colimits, we have the identification $\text{Fun}^L(\text{Sh}(X), \mathcal{C}) = \text{Fun}^L(\text{Sh}(X; \text{Sp}); \mathcal{C})$. We then consider the chain of equivalences

$$\begin{aligned} \text{Sh}(X; \mathcal{C}) &= \text{CSh}(X; \mathcal{C}) \xrightarrow{\sim} \text{Fun}^L(\text{Sh}(X, \text{Sp}), \mathcal{C}) \xleftarrow{\sim} \text{Fun}^L(\text{Sh}(X; \text{Sp}), \text{Sp}) \otimes \mathcal{C} \\ &\xleftarrow{\sim} \text{CSh}(X; \text{Sp}) \otimes \mathcal{C} = \text{Sh}(X; \text{Sp}) \otimes \mathcal{C}, \end{aligned}$$

where $\text{Sh}(X; \mathcal{C}) = \text{CSh}(X; \mathcal{C})$ is by Verdier duality (Theorem 2.6.6), and $\text{CSh}(X; \mathcal{C}) \xrightarrow{\sim} \text{Fun}^L(\text{Sh}(X, \text{Sp}), \mathcal{C})$ is the abstract identification of Lemma 2.5.11. The map bridging the scalar extension,

$$\text{Fun}^L(\text{Sh}(X, \text{Sp}), \text{Sp}) \otimes \mathcal{C} \rightarrow \text{Fun}^L(\text{Sh}(X, \text{Sp}), \mathcal{C}),$$

is a special case of the canonical abstract map, in a symmetric monoidal category $\mathcal{C}$, of the form

$$(11) \quad \underline{\text{Hom}}(x, 1_{\mathcal{C}}) \otimes y \rightarrow \underline{\text{Hom}}(x, y),$$

induced from the counit $x \otimes \underline{\text{Hom}}(x, 1_{\mathcal{C}}) \rightarrow 1_{\mathcal{C}}$ by tensoring with y and passing from $x \otimes (-)$ to $\underline{\text{Hom}}(x, -)$. This map is an equivalence when x is dualizable, which in our case is guaranteed by Corollary 2.8.15. $\square$

Notation 2.9.5. When emphasizing the coefficient $\mathcal{C}$, we will use notations such as $f_{\mathcal{C}}^*$.

The main proposition of this section is the existence of the $*$ -pullback for general coefficients.

Proposition 2.9.6 ([66, Corollary 5.22]). *Let $f : X \rightarrow Y$ be a map in LCH and let $\mathcal{C}$ be a stable category with both limits and colimits. Then the pushforward*

$$f_{\mathcal{C},*} : \mathrm{Sh}(X; \mathcal{C}) \rightarrow \mathrm{Sh}(Y; \mathcal{C})$$

admits a left adjoint

$$f_{\mathcal{C}}^* : \mathrm{Sh}(Y; \mathcal{C}) \rightarrow \mathrm{Sh}(X; \mathcal{C}).$$

Furthermore, under the identification $\mathrm{Sh}(X; \mathcal{C}) = \mathrm{Sh}(X; \mathrm{Sp}) \otimes \mathcal{C}$, we have $f_{\mathcal{C}}^ = f_{\mathrm{Sp}}^* \otimes \mathcal{C}$.*

Lemma 2.9.7. *Let $F : \mathcal{C} \rightarrow \mathcal{D}$ and $G : \mathcal{D} \rightarrow \mathcal{E}$ be functors. Assume that F^L and G^L exist. Then $(G \circ F)^L$ exists and is given by $F^L \circ G^L$.*

Proof. One checks directly that $F^L \circ G^L$ satisfies the universal property of $(G \circ F)^L$. $\square$

For any $f : X \rightarrow Y$, we can use the decomposition result of Lemma 2.8.4 to factorize $f = p \circ j$ for some open embedding j and proper map p . Since $f_{\mathcal{C},*} = p_{\mathcal{C},*} \circ j_{\mathcal{C},*}$, the above lemma implies that it suffices to prove the cases when $f = p$ is proper and when $f = j$ is an open embedding separately. We would like to trace what is known through the chain of maps giving the identification $\mathrm{Sh}(X; \mathcal{C}) = \mathrm{Sh}(X; \mathrm{Sp}) \otimes \mathcal{C}$. The passage between sheaves and cosheaves is compared in Lemmas 2.6.10 and 2.6.11. The passage

$$\mathrm{Fun}^L(\mathrm{Sh}(X; \mathrm{Sp}), \mathrm{Sp}) \otimes \mathcal{C} \rightarrow \mathrm{Fun}^L(\mathrm{Sh}(X; \mathrm{Sp}), \mathcal{C})$$

is more abstract, and the following lemma can be read off directly from the discussion at the end of the proof of Lemma 2.9.4.

Lemma 2.9.8. *Let $\mathcal{C}$ be a closed symmetric monoidal category. For a morphism $\alpha : x_1 \rightarrow x_2$ in $\mathcal{C}$, we denote by*

$$f^{\vee, y} : \underline{\mathrm{Hom}}(x_2, y) \rightarrow \underline{\mathrm{Hom}}(x_1, y)$$

the induced morphism under the anti-endomorphism $\underline{\mathrm{Hom}}(-, y) : \mathcal{C}^{\mathrm{op}} \rightarrow \mathcal{C}$. Then we have the following commutative diagram

$$\begin{array}{ccc} \underline{\mathrm{Hom}}(x_2, 1_{\mathcal{C}}) \otimes y & \longrightarrow & \underline{\mathrm{Hom}}(x_2, y) \\ \downarrow f^{\vee, 1_{\mathcal{C}}} \otimes \mathrm{id}_y & & \downarrow f^{\vee, y} \\ \underline{\mathrm{Hom}}(x_1, 1_{\mathcal{C}}) \otimes y & \longrightarrow & \underline{\mathrm{Hom}}(x_1, y) \end{array}$$

where the horizontal maps are the maps given in (11).

Lastly, we identify the passage between $\mathrm{CSh}(X; \mathcal{C})$ and $\mathrm{Fun}^L(\mathrm{Sh}(X; \mathcal{C})) = \mathrm{Fun}^L(\mathrm{Sh}(X; \mathrm{Sp}); \mathcal{C})$.

Lemma 2.9.9. *Let $j : U \rightarrow X$ be an open embedding. Then $j^* : \mathrm{Sh}(X) \rightarrow \mathrm{Sh}(U)$ admits a left adjoint*

$$j_{\#} : \mathrm{Sh}(U) \rightarrow \mathrm{Sh}(X).$$

Proof. Since j^* is given simply by restricting the testing open sets, we see that, if $\{F_{\alpha}\}$ is a diagram of sheaves on U indexed by I , the restrictions $j^*(\lim_{\alpha} F_{\alpha})$ and $j^*(\mathrm{colim}_{\alpha} F_{\alpha})$ tautologically satisfy the required mapping properties. We conclude that j^* preserves all small limits and colimits (and in particular κ -filtered colimits for any κ). By the adjoint

functor theorem (Theorem 2.2.6), the left adjoint $j_\#$ exists. Alternatively, one can verify that, for $G \in \text{Sh}(U)$, the functor $j_\#$ is concretely given by

$$(j_\# G)(V) = \begin{cases} G(V), & V \subseteq U, \\ \emptyset, & \text{otherwise.} \end{cases}$$

□

Lemma 2.9.10. *Let $j : U \rightarrow X$ be an open embedding and let $\mathcal{C}$ be stable with limits and colimits. The cosheaf pullback*

$$j^{!,\text{co}} : \text{CSh}(X; \mathcal{C}) \rightarrow \text{CSh}(U; \mathcal{C})$$

corresponds to

$$(-) \circ j_\# : \text{Fun}^L(\text{Sh}(X), \mathcal{C}) \rightarrow \text{Fun}^L(\text{Sh}(U), \mathcal{C})$$

under the identification $\text{Fun}^L(\text{Sh}(X); \mathcal{C}) = \text{CSh}(X; \mathcal{C})$.

Proof. The adjunction $j_\# \dashv j^*$ is internal in $\widehat{\text{Cat}}^{\text{colim}}$, so $\text{Fun}^L(-, \mathcal{C})$, which extends naturally to a 2-functor, sends it to the adjunction $(-) \circ j^* \dashv (-) \circ j_\#$. We can then conclude by Lemma 2.9.1. Alternatively, using the concrete description of $j_\#$, we can compute, similarly to Lemma 2.9.1, that for $G \in \text{Fun}^L(\text{Sh}(X); \mathcal{C})$ and $V \subseteq U$,

$$(j^{!,\text{co}} \bar{G})(V) = \bar{G}(V) = G(h_V),$$

and similarly

$$(\overline{G \circ j_\#})(V) = G(j_\# h_V) = G(h_V).$$

□

Proof of Proposition 2.9.6. Decompose $f = p \circ j$ for some open embedding j and proper map p . Since $f_{\mathcal{C},*} = p_{\mathcal{C},*} \circ j_{\mathcal{C},*}$, it suffices to prove the statement for $p_{\mathcal{C},*}$ and $j_{\mathcal{C},*}$ separately. Under the chain of identifications recalled in Lemma 2.9.4, the functor $j_{\mathcal{C}}^*$ is identified via

$$j_{\mathcal{C}}^* \leftrightarrow j_{\mathcal{C}}^{!,\text{co}} \leftrightarrow (-)_{\mathcal{C}} \circ j^* \leftrightarrow ((-)_{\text{Sp}} \circ j^*) \otimes \mathcal{C} \leftrightarrow j_{\text{Sp}}^{!,\text{co}} \otimes \mathcal{C} \leftrightarrow j_{\text{Sp}}^* \otimes \mathcal{C},$$

where we use the notation $(-)_{\mathcal{C}} \circ$ to indicate the target coefficient for $\text{Fun}^L(-, \mathcal{C})$, by Lemmas 2.6.10, 2.9.8, and 2.9.1. Similarly, $p_{\mathcal{C},*}$ is identified via

$$p_{\mathcal{C},*} \leftrightarrow p_{\mathcal{C},!}^{\text{co}} \leftrightarrow (-)_{\mathcal{C}} \circ p^* \leftrightarrow ((-)_{\text{Sp}} \circ p^*) \otimes \mathcal{C} \leftrightarrow p_{\text{Sp},!}^{\text{co}} \otimes \mathcal{C} \leftrightarrow p_{\text{Sp},*} \otimes \mathcal{C}$$

by Lemmas 2.6.11, 2.9.8, and 2.9.10. Passing to the left adjoint, we see that $p_{\mathcal{C}}^* := p_{\mathcal{C},*}^L$ corresponds to $p_{\text{Sp}}^* \otimes \mathcal{C}$, so

$$f_{\mathcal{C}}^* \leftrightarrow (j_{\text{Sp}}^* \otimes \mathcal{C}) \circ (p_{\text{Sp}}^* \otimes \mathcal{C}) = (j_{\text{Sp}}^* \circ p_{\text{Sp}}^*) \otimes \mathcal{C} = f_{\text{Sp}}^* \otimes \mathcal{C}.$$

□

Remark 2.9.11. Note that we did not use the existence of the sheafification functor to show that $f_{\mathcal{C}}^*$ exists in general. In fact, the proof of the existence of sheafification given in [66, Theorem 5.23] uses a similar scalar extension argument.

As a corollary, we have a general version of the Künneth formula.

Corollary 2.9.12. *Let X and Y be locally compact Hausdorff spaces and let $\mathcal{C}$ and $\mathcal{D}$ be stable categories with limits and colimits. Then we have*

$$\mathrm{Sh}(X; \mathcal{C}) \otimes \mathrm{Sh}(Y; \mathcal{D}) = \mathrm{Sh}(X \times Y; \mathcal{C} \otimes \mathcal{D}).$$

Furthermore, if $f : X' \rightarrow X$ and $g : Y' \rightarrow Y$ are continuous maps, then the following diagram commutes

$$\begin{array}{ccc} \mathrm{Sh}(X; \mathcal{C}) \otimes \mathrm{Sh}(Y; \mathcal{D}) & \longrightarrow & \mathrm{Sh}(X \times Y; \mathcal{C} \otimes \mathcal{D}) \\ \downarrow f_{\mathcal{C}}^* \otimes g_{\mathcal{D}}^* & & \downarrow (f \times g)_{\mathcal{C} \otimes \mathcal{D}}^* \\ \mathrm{Sh}(X'; \mathcal{C}) \otimes \mathrm{Sh}(Y'; \mathcal{D}) & \longrightarrow & \mathrm{Sh}(X' \times Y'; \mathcal{C} \otimes \mathcal{D}). \end{array}$$

We note that a similar argument for cosheaves implies that, by going only halfway through the argument, we have $f_{\mathcal{C},!}^{\mathrm{co}} \leftrightarrow f_{\mathrm{Sp},!}^{\mathrm{co}} \otimes \mathcal{C}$ under the identification $\mathrm{CSh}(X; \mathcal{C}) = \mathrm{CSh}(X; \mathrm{Sp}) \otimes \mathcal{C}$. Here, we implicitly use Remark 2.8.16 to show that $f_{\mathrm{Sp},!}^{\mathrm{co}}$ is a morphism in Pr^{L} . This in turn implies that $f_{\mathcal{C},!}^{\mathrm{co}}$ is also a morphism in Pr^{L} and thus admits an underlying right adjoint $f_{\mathcal{C}}^{\mathrm{!},\mathrm{co}}$. Therefore,

Proposition 2.9.13. *Assumption 2.7.1 holds for any stable category $\mathcal{C}$ with small limits and colimits. That is, if $p : X \rightarrow Y$ is proper, then $p_{\mathcal{C},*}$ has a right adjoint $p_{\mathcal{C}}^!$ and, if $j : U \rightarrow U$ is open, then $j_{\mathcal{C}}^*$ has a left adjoint $j_{\mathcal{C},!}$.*

2.10. Six-functor formalism in topology: General coefficients. With Proposition 2.9.13 established, we can proceed with our discussion of obtaining the six-functor formalism for general coefficients. However, to state the projection formula, we need to take a detour and discuss the notion of symmetric monoidal stable presentable categories.

Approximate Definition 2.10.1. Let $(\mathcal{C}, \otimes)$ be a symmetric monoidal category. A commutative algebra object A in $(\mathcal{C}, \otimes)$, often simply written as $A \in \mathrm{CAlg}(\mathcal{C})$, is an object $A \in \mathcal{C}$ equipped with a multiplication $m : A \otimes A \rightarrow A$ and a unit map $i : 1_{\mathcal{C}} \rightarrow A$, satisfying the usual requirements of a commutative algebra together with higher coherences.

Definition 2.10.2. A symmetric monoidal stable presentable category $\mathcal{V}$ is an object $\mathcal{V} \in \mathrm{CAlg}(\mathrm{Pr}_{st}^{\mathrm{L}})$, where $\mathrm{Pr}_{st}^{\mathrm{L}}$ is equipped with the Lurie tensor product (Notation 2.5.4).

Concretely, $\mathcal{V} \in \mathrm{CAlg}(\mathrm{Pr}^{\mathrm{L}})$ is a presentable category equipped with a symmetric monoidal structure $(\mathcal{V}, \otimes_{\mathcal{V}}, 1_{\mathcal{V}})$ such that the assignment

$$(M, N) \mapsto M \otimes_{\mathcal{V}} N$$

is colimit-preserving in both M and N . Indeed, the multiplication $\mathcal{V} \otimes \mathcal{V} \rightarrow \mathcal{V}$ supplies the tensor product $\otimes_{\mathcal{V}}$, the unit $\mathrm{Sp} \rightarrow \mathcal{V}$ supplies the tensor unit, and the higher coherences are dictated by those of a commutative algebra object.

Example 2.10.3. Part (4) of Proposition 2.5.5 implies that Sp , being the unit of $\mathrm{Pr}_{st}^{\mathrm{L}}$, is in particular a commutative algebra object. The symmetric monoidal structure on Sp can be concretely realized by the smash product (see, for example, [45, Section 4.8.2] for details) from algebraic topology, with unit given by $\mathbb{S} = \Sigma_{\mathrm{Ani},+}^{\infty}(*)$, the sphere spectrum, colloquially called “the sphere.” The name comes from the classical construction of $\mathbb{S}$: one first considers the suspension pre-spectrum over the point $*$, which can be concretely described by the sequence of n -spheres $(\cdots, S^2, S^1, S^0)$ viewed as pointed anima. This is a pre-spectrum

since the canonical map $S^n \rightarrow \Omega_* S^{n+1}$ is not an equivalence, but one can spectrify it to $\mathbb{S}$ by setting the n -th term to be $\operatorname{colim}_{k \rightarrow \infty} \Omega_*^k S^{n+k}$.

Remark 2.10.4. Iterating the discussion of Example 2.10.3, since $\mathbb{S}$ is the unit in Sp , it is in particular the initial object in $\operatorname{CAlg}(\operatorname{Sp})$, i.e., the initial commutative ring spectrum. There is a notion of connective spectra $\operatorname{Sp}^{\operatorname{cn}} \subseteq \operatorname{Sp}$, consisting of those $X \in \operatorname{Sp}$ with $\pi_n(X) = 0$ for $n < 0$,²³ and the symmetric monoidal structure restricts to $\operatorname{Sp}^{\operatorname{cn}}$, so $\mathbb{S}$ is the unit there as well. Inside $\operatorname{Sp}^{\operatorname{cn}}$, the functor $\pi_0 : \operatorname{Sp}^{\operatorname{cn}} \rightarrow \operatorname{Ab} \subseteq \operatorname{Sp}^{\operatorname{cn}}$ is left exact and sends commutative algebras to commutative algebras, yielding a canonical map $\mathbb{S} \rightarrow \pi_0(\mathbb{S}) = \mathbb{Z}$. This is the precise sense in which one says “the sphere is the absolute base ring.”

Approximate Definition 2.10.5. Let $(\mathcal{C}, \otimes)$ be a symmetric monoidal category and let $A \in \operatorname{CAlg}(\mathcal{C})$. An A -module in $\mathcal{C}$ consists of an object $X \in \mathcal{C}$ together with an action map $a : A \otimes X \rightarrow X$, subject to the usual module compatibility conditions together with higher coherences. The category of A -modules is denoted $\operatorname{Mod}_A(\mathcal{C})$.

Example 2.10.6. Let X be a topological space. Then $\operatorname{Sh}(X; \operatorname{Sp})$ is a Sp -module in Pr_{st}^L . Concretely, the action $\operatorname{Sp} \otimes \operatorname{Sh}(X; \operatorname{Sp}) \rightarrow \operatorname{Sh}(X; \operatorname{Sp})$ is given by the assignment

$$(A, F) \mapsto A_X \otimes F,$$

for $F \in \operatorname{Sh}(X; \operatorname{Sp})$ and $A \in \operatorname{Sp}$, where $A_X := a^*A$ and $a : X \rightarrow *$ is the projection to a point.

Definition 2.10.7. Let $(\mathcal{C}, \otimes)$ be a symmetric monoidal category and let $A \in \operatorname{CAlg}(\mathcal{C})$ be a commutative algebra object. For two A -modules $X, Y \in \operatorname{Mod}_A(\mathcal{C})$, their relative tensor product over A , denoted $X \otimes_A Y$, is an A -module object equipped with a canonical map $X \times Y \rightarrow X \otimes_A Y$ such that, for any $Z \in \operatorname{Mod}_A(\mathcal{C})$, the induced pre-composition map

$$\operatorname{Hom}_{\operatorname{Mod}_A(\mathcal{C})}(X \otimes_A Y, Z) \hookrightarrow \operatorname{Hom}_{\mathcal{C}}(X \times Y, Z)$$

is a monomorphism whose image consists of maps $X \times Y \rightarrow Z$ that are A -bilinear.

We refer the readers to [45, Theorem 4.5.3.1] for the precise version of the next two statements.

Lemma 2.10.8. *With the same setting as Definition 2.10.7, if $\mathcal{C}$ admits colimits, then the relative tensor product $\otimes_A$ exists and canonically equips $\operatorname{Mod}_A(\mathcal{C})$ with a symmetric monoidal structure, with A as its unit.*

Proof. For $X, Y \in \operatorname{Mod}_A(\mathcal{C})$, the tensor product $X \otimes_A Y$ is given by the geometric realization of the cosimplicial object $(n \mapsto X \otimes A^{\otimes n} \otimes Y)$, i.e., we have

$$\operatorname{colim}_{\Delta} \left(\cdots \rightrightarrows X \otimes A \otimes A \otimes Y \rightrightarrows X \otimes A \otimes Y \rightrightarrows X \otimes Y \right) \xrightarrow{\sim} X \otimes_A Y.$$

□

Proposition 2.10.9. *Let $(\mathcal{C}, \otimes)$ be a symmetric monoidal category with small colimits and let $A \in \operatorname{CAlg}(\mathcal{C})$ be a commutative algebra object. Then the forgetful functor*

$$\operatorname{Mod}_A(\mathcal{C}) \rightarrow \mathcal{C}$$

admits a symmetric monoidal left adjoint given by the free module functor

$$\begin{aligned} \mathcal{C} &\rightarrow \operatorname{Mod}_A(\mathcal{C}) \\ X_0 &\mapsto A \otimes X_0. \end{aligned}$$

²³See the discussion in Example 2.4.13.

Notation 2.10.10. Let $\mathcal{V}$ be a stable symmetric monoidal presentable category. We denote $\mathrm{Pr}_{\mathcal{V}}^{\mathrm{L}} := \mathrm{Mod}_{\mathcal{V}}(\mathrm{Pr}^{\mathrm{L}})$ and call it the category of $\mathcal{V}$ -linear presentable categories.

Example 2.10.11. From this point of view, part (4) of Proposition 2.5.5 implies that Sp -linear presentable categories $\mathrm{Pr}_{\mathrm{Sp}}^{\mathrm{L}}$ are simply stable presentable categories $\mathrm{Pr}_{st}^{\mathrm{L}}$. By Remark 2.10.4, any classical commutative ring R (sometimes called a static commutative ring) can be viewed as a connective ring spectrum. More generally, applying Lemma 2.10.8 one level down, for any ring spectrum R we can consider the presentable symmetric monoidal category $\mathrm{Mod}_R := \mathrm{Mod}_R(\mathrm{Sp})$. In this case, we use the notation $\mathrm{Pr}_R^{\mathrm{L}} := \mathrm{Pr}_{\mathrm{Mod}_R}^{\mathrm{L}}$ and call these R -linear presentable categories. If R is a classical commutative ring, this category contains the classical notion of pretriangulated differential graded categories over R ; see [17] for details.

Fix a symmetric monoidal stable presentable category $\mathcal{V}$. We will use the notation $\otimes_{\mathcal{V}}$ to denote the symmetric monoidal product in $\mathrm{Pr}_{\mathcal{V}}^{\mathrm{L}}$. We will establish the theorem below, which gives the six-functor formalism in topology over $\mathcal{V}$, through several lemmas.

Theorem 2.10.12. *Let $\mathcal{V}$ be a symmetric monoidal stable presentable category.²⁴ The functor*

$$\begin{aligned} \mathrm{Sh}^* : \mathrm{LCH}^{\mathrm{op}} &\rightarrow \mathrm{Pr}_{\mathcal{V}}^{\mathrm{L}} \\ X &\mapsto \mathrm{Sh}(X; \mathcal{V}) \end{aligned}$$

is symmetric monoidal. The families P and I given by proper maps and open immersions satisfy the conditions in Proposition 2.8.5. As a result, it extends to a six-functor formalism $\mathrm{Sh} : \mathrm{Corr}(\mathrm{LCH})^{\otimes} \rightarrow \mathrm{Pr}_{\mathcal{V}}^{\mathrm{L}}$.

That Sh^* is symmetric monoidal is a special case of Corollary 2.9.12, and is concretely stated as follows:

Lemma 2.10.13. *Let X and Y be locally compact Hausdorff spaces. Then we have*

$$\mathrm{Sh}(X; \mathcal{V}) \otimes_{\mathcal{V}} \mathrm{Sh}(Y; \mathcal{V}) = \mathrm{Sh}(X \times Y; \mathcal{V}).$$

Furthermore, if $f : X' \rightarrow X$ and $g : Y' \rightarrow Y$ are continuous maps, then the following diagram commutes

$$\begin{array}{ccc} \mathrm{Sh}(X; \mathcal{V}) \otimes_{\mathcal{V}} \mathrm{Sh}(Y; \mathcal{V}) & \longrightarrow & \mathrm{Sh}(X \times Y; \mathcal{V}) \\ \downarrow f^* \otimes_{\mathcal{V}} g^* & & \downarrow (f \times g)^* \\ \mathrm{Sh}(X'; \mathcal{V}) \otimes_{\mathcal{V}} \mathrm{Sh}(Y'; \mathcal{V}) & \longrightarrow & \mathrm{Sh}(X' \times Y'; \mathcal{V}). \end{array}$$

Definition 2.10.14. The symmetric monoidal product $\otimes$ on $\mathrm{Sh}(X; \mathcal{V})$ is defined to be the composition

$$\begin{aligned} \mathrm{Sh}(X; \mathcal{V}) \otimes_{\mathcal{V}} \mathrm{Sh}(X; \mathcal{V}) &= (\mathrm{Sh}(X; \mathrm{Sp}) \otimes \mathcal{V}) \otimes_{\mathcal{V}} (\mathrm{Sh}(X; \mathrm{Sp}) \otimes \mathcal{V}) \\ &= (\mathrm{Sh}(X; \mathrm{Sp}) \otimes \mathrm{Sh}(X; \mathrm{Sp})) \otimes \mathcal{V} \\ &= \mathrm{Sh}(X \times X; \mathrm{Sp}) \otimes \mathcal{V} = \mathrm{Sh}(X \times X; \mathcal{V}) \xrightarrow{\Delta^*} \mathrm{Sh}(X; \mathcal{V}). \end{aligned}$$

Specializing Lemma 2.10.13 to the case $Y = *$, we conclude:

²⁴As shown in [66, Section 6], the same argument works more generally for symmetric monoidal stable categories with both limits and colimits such that the tensor structure is stable and colimit-preserving in each component.

Lemma 2.10.15. *Let X be a locally compact Hausdorff space. Then $\mathrm{Sh}(X; \mathcal{V})$ is a $\mathcal{V}$ -module in Pr^{L} . Concretely, denote by $a : X \rightarrow *$ the projection to a point. The $\mathcal{V}$ -module action is given by*

$$\begin{aligned} \mathcal{V} \otimes \mathrm{Sh}(X; \mathcal{V}) &\rightarrow \mathrm{Sh}(X; \mathcal{V}) \\ A \otimes F &\mapsto A_X \otimes F, \end{aligned}$$

where $A_X := a^*A$. For any $f : X \rightarrow Y$ in LCH , the pullback $f^* : \mathrm{Sh}(Y; \mathcal{V}) \rightarrow \mathrm{Sh}(X; \mathcal{V})$ is symmetric monoidal and, in particular, $\mathcal{V}$ -linear.

As mentioned, Lemmas 2.8.6 and 2.8.7 do not depend on the coefficient. To obtain proper base change, we need to again perform an extension of scalars, using Lemma 2.8.8.

Lemma 2.10.16. *Let $\mathcal{V}$ be a symmetric monoidal stable presentable category. The functor Sh^* satisfies base change for proper maps. That is, if $p : X \rightarrow Y$ is proper and*

$$\begin{array}{ccc} X' & \xrightarrow{q'} & X \\ \downarrow p' & & \downarrow p \\ Y' & \xrightarrow{q} & Y \end{array}$$

*is a pullback diagram, then the canonical map $q^*p_* \rightarrow p'_*q'^*$ is an equivalence.*

Proof. By Corollary 2.9.6, for any $f : X \rightarrow Y$ in LCH , we have $f_{\mathcal{V}}^* = f_{\mathrm{Sp}}^* \otimes \mathcal{V}$ under the identification $\mathrm{Sh}(X; \mathcal{V}) = \mathrm{Sh}(X; \mathrm{Sp}) \otimes \mathcal{V}$. By Corollary 2.9.2, if $p : X \rightarrow Y$ is proper, then $p_{\mathrm{Sp},*}$ admits a right adjoint, so $p_{\mathrm{Sp}}^* \dashv p_{\mathrm{Sp},*}$ is an adjunction internal to Pr^{L} . Tensoring with $\mathcal{V}$, we see that $p_{\mathcal{V},*}$, the underlying right adjoint of $p_{\mathcal{V}}^*$, is $\mathcal{V}$ -linear and colimit-preserving, and in particular corresponds to $p_{\mathrm{Sp},*} \otimes \mathcal{V}$. Tensoring the equivalence $q^*p_* \xrightarrow{\sim} p'_*q'^*$ over the sphere with $\mathcal{V}$ then yields the same equivalence over $\mathcal{V}$. $\square$

We are again left with the projection formulas.

Lemma 2.10.17. *Let $f : X \rightarrow Y$ be a map in LCH , $F \in \mathrm{Sh}(X; \mathcal{V})$, and $G \in \mathrm{Sh}(Y; \mathcal{V})$. We have the following projection formulas:*

(1) *Assume $f = j$ is an open embedding. Then there is a canonical map*

$$j_!(F \otimes j^*G) \rightarrow (j_!F) \otimes G,$$

which is an equivalence.

(2) *Assume $f = p$ is proper. Then there is a canonical map*

$$(p_*F) \otimes G \rightarrow p_*(F \otimes p^*G),$$

which is an equivalence.

Proof. The proof proceeds in the same way as Lemma 2.8.13, once all the relevant functors are obtained via extension of scalars. Alternatively, the equivalences over $\mathcal{V}$ can be obtained from those over the sphere by extension of scalars. $\square$

Remark 2.10.18. For a general map $f : X \rightarrow Y$ in LCH , one can likewise deduce the projection formula from base change of $!$ -pushforward against $*$ -pullback. However, like base change itself, the isomorphism

$$(f_!F) \otimes G = f_!(F \otimes f^*G)$$

is a piece of data rather than merely a property.

Remark 2.10.19. A more abstract way to phrase the projection formula is as follows. For a map $f : X \rightarrow Y$, we noted in Lemma 2.10.15 that $f^* : \mathrm{Sh}(Y; \mathcal{V}) \rightarrow \mathrm{Sh}(X; \mathcal{V})$ is symmetric monoidal. This means that f^* is a $\mathcal{V}$ -algebra map, i.e., f^* is a morphism in $\mathrm{CAlg}(\mathrm{Pr}_{\mathcal{V}}^{\mathrm{L}})$, so $\mathrm{Sh}(X; \mathcal{V})$ acquires the structure of a $\mathrm{Sh}(Y; \mathcal{V})$ -module. The projection formula is then equivalent to the statement that the $\mathcal{V}$ -linear map

$$f_! : \mathrm{Sh}(X; \mathcal{V}) \rightarrow \mathrm{Sh}(Y; \mathcal{V})$$

is in fact a $\mathrm{Sh}(Y; \mathcal{V})$ -module map.

3. MICROLOCAL SHEAF THEORY

Warning 3.0.1. As we are about to complete the foundational aspects of microlocal sheaf theory, we will soon restrict ourselves to the case when the coefficient category $\mathcal{V}$ is a symmetric monoidal stable ω -accessible presentable category and work exclusively with $\mathcal{V}$ -valued sheaves, more precisely after Assumption 3.2.17. As a result, unless otherwise specified, we will adopt the simplified notation $\mathrm{Sh}(M) := \mathrm{Sh}(M; \mathcal{V})$ after that point, as anima-valued sheaves will seldom be considered again.

3.1. Stalks and support. When working with sheaves on a topological space X , one often reaches a situation where there are two sheaves $F, G \in \mathrm{Sh}(X; \mathcal{C})$ and a morphism $\alpha : F \rightarrow G$ that one would like to verify is an equivalence. In the most general situation, the best one can hope for is that the equivalence is easy to see “locally”: that there is a large enough class of open sets U such that $\alpha|_U : F|_U \xrightarrow{\sim} G|_U$ is an equivalence, and then one uses the sheaf condition (1) to conclude that α itself is an equivalence.

Classically, when $\mathcal{C}$ is the category of abelian groups, the situation can be further simplified as equivalences can be checked pointwise. That is, if $F \in \mathrm{Sh}(X; \mathrm{Ab})$ is a sheaf valued in abelian groups, one can define the stalk $F_x \in \mathrm{Ab}$ as the germs at x : a section $s \in F_x$ is an equivalence class $[(s_U, U)]$, where $U \ni x$ is an open neighborhood of x and $s_U \in F(U)$ is a section over U . Two such pairs (s_1, U_1) and (s_2, U_2) are equivalent if and only if there exists an open set $W \ni x$ such that $W \subseteq U_1 \cap U_2$ and $s_1|_W = s_2|_W$. The key fact is that a morphism $\alpha : F \rightarrow G$ is an equivalence if and only if $\alpha_x : F_x \rightarrow G_x$ is an equivalence at every stalk. We note that in the higher categorical setting this is not true in general, but there are reasonable finiteness conditions one can impose, both on the coefficient category and on the topological space, that make the statement still hold.

Definition 3.1.1. Let $\mathcal{C}$ be a category with filtered colimits, let X be a topological space, and let $x \in X$ be a point. We denote by $i_x : * \hookrightarrow X$ the inclusion of the point x . The stalk functor at x is the functor

$$i_x^* : \mathrm{Sh}(X; \mathcal{C}) \rightarrow \mathrm{Sh}(*; \mathcal{C}) = \mathcal{C},$$

and for $F \in \mathrm{Sh}(X; \mathcal{C})$, we write $F_x := i_x^* F$ and call it the stalk of F at x . Concretely, it is given by

$$F_x = \mathrm{colim}_{U \ni x} F(U),$$

where the colimit runs over all open neighborhoods of x .

Without going into too much detail, we give the finiteness condition on the space:

Theorem 3.1.2 ([44, Corollary 7.2.1.17]). *Let X be a topological space such that $\mathrm{Sh}(X; \mathrm{Ani})$ is of finite homotopical dimension, e.g., a finite-dimensional manifold. Then $\mathrm{Sh}(X; \mathrm{Ani})$ has*

enough points. That is, a morphism $\alpha : F \rightarrow G$ in $\mathrm{Sh}(X; \mathrm{Ani})$ is an equivalence if and only if $\alpha_x : F_x \rightarrow G_x$ is an equivalence for all $x \in X$.

Remark 3.1.3. The above theorem is a consequence of several facts from higher topos theory:

- (1) Recall that Whitehead's theorem asserts that a map $f : X \rightarrow Y$ between anima is an equivalence if and only if the induced map $\pi_0(X) \rightarrow \pi_0(Y)$ is an equivalence and, for any $x_0 \in \pi_0(X)$, the induced map $\pi_n(X, x_0) \rightarrow \pi_n(Y, f(x_0))$ is an equivalence for all n . A notion of homotopy groups can be defined for any topos $\mathcal{X}$ [44, Definition 6.5.1.1], and maps satisfying the latter conditions are called ∞ -connective [44, Definition 6.5.1.10].
- (2) A topos $\mathcal{X}$ is called hypercomplete if ∞ -connective maps are invertible, and any topos admits a hypercompletion $L : \mathcal{X} \rightarrow \widehat{\mathcal{X}}$, which is a Bousfield localization [44, Lemma 6.5.2.10] (and the paragraph after it). The right adjoint $\widehat{\mathcal{X}} \subseteq \mathcal{X}$ realizes it as the subcategory consisting of hypercomplete objects, i.e., those $X \in \mathcal{X}$ such that $X \rightarrow *$ is ∞ -connective. In the case $\mathcal{X} = \mathrm{Sh}(X; \mathrm{Ani})$, the subcategory $\widehat{\mathcal{X}}$ is given by $\widehat{\mathrm{Sh}}(X; \mathrm{Ani})$, the subcategory of sheaves satisfying the stronger descent condition along arbitrary hypercovers [44, Theorem 6.5.3.12].
- (3) Let X be a topological space. While having $\alpha_x : F_x \rightarrow G_x$ be an equivalence for all x does not in general imply that a morphism $\alpha : F \rightarrow G$ in $\mathrm{Sh}(X; \mathrm{Ani})$ is an equivalence, one can nevertheless localize at such morphisms. As explained in [44, Item (6) of Section 6.5.4], this localization is exactly $\widehat{\mathrm{Sh}}(X; \mathrm{Ani}) \subseteq \mathrm{Sh}(X; \mathrm{Ani})$.
- (4) One then proves that locally having finite homotopical dimension implies hypercompleteness [44, Corollary 7.2.1.12]. In this case $\widehat{\mathrm{Sh}}(X; \mathrm{Ani}) = \mathrm{Sh}(X; \mathrm{Ani})$, and equivalences can be detected at stalks. More directly, as argued in [44, Corollary 7.2.1.17], one begins with the fact that, under the finite-dimensional assumption, $\mathrm{Sh}(X; \mathrm{Ani})$ is hypercomplete, and then argues that being an equivalence on all stalks implies ∞ -connectivity.
- (5) To obtain the statement for manifolds, recall that an n -dimensional manifold has covering dimension $\leq n$. Then by [44, Theorem 7.2.3.6], a paracompact topological space X has $\mathrm{Sh}(X; \mathrm{Ani})$ of homotopical dimension $\leq n$.

Convention 3.1.4. We will simply say that X is finite-dimensional to mean that it is paracompact and of finite covering dimension.

We now turn our attention to the coefficient category. We begin with the theory of idempotent completeness [47, Tag 03Y9].

Definition 3.1.5. Let $\mathcal{C}$ be a category.²⁵

- (1) A map $r : X \rightarrow Y$ is a retract if there exists a map $\iota : Y \hookrightarrow X$ such that $r \circ \iota = \mathrm{id}_Y$.
- (2) A morphism $e : X \rightarrow X$ in $\mathcal{C}$ is called an idempotent if $e \circ e = e$. Note that any retract induces an idempotent via $e = \iota \circ r$.
- (3) The category $\mathcal{C}$ is said to be idempotent complete if every idempotent e arises from a retract.

Lemma 3.1.6 ([47, Tag 040N]). *Let $\mathcal{C}$ be a category with sequential colimits (or limits). Then $\mathcal{C}$ is idempotent complete.*

²⁵See [47, Tag 03ZT], [47, Tag 03ZX], and [47, Tag 0406] for more precise definitions.

Proof. Let $e : X \rightarrow X$ be an idempotent. The colimit (or limit) of the sequence

$$\dots \xrightarrow{e} X \xrightarrow{e} X \xrightarrow{e} X \xrightarrow{e} \dots$$

provides a retract for e . □

Ultimately, we are concerned with the dictionary between large and small categories, and idempotent completeness is a condition needed to ensure a one-to-one correspondence.

Example 3.1.7. A higher categorical version of Example 2.2.16 is the following. Denote by $\mathbb{Z}$ the ring of integers, viewed as a ring spectrum via Remark 2.10.4. Then the category of $\mathbb{Z}$ -modules, $\text{Mod}_{\mathbb{Z}}$, is the Ind-completion of the category $\{\mathbb{Z}\}$ consisting of a single object whose endomorphism ring is $\mathbb{Z}$. Denote by $\text{Perf}(\mathbb{Z})$ the full subcategory of $\text{Mod}_{\mathbb{Z}}$ consisting of those $M \in \text{Mod}_{\mathbb{Z}}$ such that $\pi_n(M) = 0$ for all but finitely many $n \in \mathbb{Z}$ and $\pi_n(M)$ is of finite rank whenever nonzero. Then the inclusion $\{\mathbb{Z}\} \subseteq \text{Perf}(\mathbb{Z})$ yields an equivalence $\text{Ind}(\{\mathbb{Z}\}) = \text{Ind}(\text{Perf}(\mathbb{Z})) = \text{Mod}_{\mathbb{Z}}$. In fact, $\text{Perf}(\mathbb{Z}) = (\text{Mod}_{\mathbb{Z}})^{\omega}$ is the subcategory of compact objects.

Notation 3.1.8. We denote by Cat^{Rex} the category of small idempotent-complete categories admitting finite colimits, with morphisms given by functors that preserve finite colimits. Similarly, we denote by $\text{Pr}_{\omega}^{\text{L}}$ the category of ω -accessible presentable categories, with morphisms given by colimit-preserving functors that preserve compact objects.

Proposition 3.1.9 ([47, Tag 067M], [44, Proposition 5.5.7.10]). *The assignment $\mathcal{C}_0 \mapsto \text{Ind}(\mathcal{C}_0)$ induces an equivalence between small idempotent-complete categories and compactly generated categories, with inverse given by $\mathcal{C} \mapsto \mathcal{C}^{\omega}$. In particular, restricting to Cat^{Rex} yields an equivalence*

$$\begin{aligned} \text{Ind} : \text{Cat}^{\text{Rex}} &\xrightarrow{\sim} \text{Pr}_{\omega}^{\text{L}} \\ \mathcal{C}_0 &\mapsto \text{Ind}(\mathcal{C}_0). \end{aligned}$$

Remark 3.1.10. There is likewise a version for a general regular cardinal $\kappa > \omega$. Since such κ would be uncountable, by Lemma 3.1.6, the corresponding category $\text{Cat}^{\text{Rex}(\kappa)}$, consisting of small categories with κ -small colimits, is automatically idempotent complete.

Now, for the purpose of studying stalks, we need a related notion of generation.

Definition 3.1.11. Let $\mathcal{C}$ be a category. A family of objects $\mathcal{S} = \{S_{\alpha}\}$ is called a generating family if, for any morphism $f : X \rightarrow Y$, f is an equivalence if and only if the induced map of anima

$$f \circ (-) : \text{Hom}_{\mathcal{C}}(S_{\alpha}, X) \rightarrow \text{Hom}_{\mathcal{C}}(S_{\alpha}, Y)$$

is an equivalence for all $S_{\alpha} \in \mathcal{S}$. When $\mathcal{S}$ consists of a single object S , we call it a generator.

Lemma 3.1.12. *Let $\mathcal{C}$ be compactly generated by a set of objects $\mathcal{S} \subseteq \mathcal{C}^{\omega}$. Then $\mathcal{S}$ forms a generating family in $\mathcal{C}$.*

Proof. The category $\mathcal{C}$ is equivalent to the Ind-completion $\text{Ind}(\mathcal{S})$, where we view $\mathcal{S}$ as a full subcategory of $\mathcal{C}$. Then Remark 2.2.15 implies that we can view $\text{Ind}(\mathcal{S}) \subseteq \text{PSh}(\mathcal{S})$ as a full subcategory of the presheaf category. That $\mathcal{S}$ is a generating family in $\mathcal{C}$ then follows, since a morphism $f : F \rightarrow G$ in $\text{PSh}(\mathcal{S})$ can be tested on objects and $\text{Hom}(h_S, F) = F(S)$ for $S \in \mathcal{S}$ by the Yoneda lemma. □

Proposition 3.1.13. *Let $\mathcal{C}$ be an ω -accessible presentable category and let X be of finite dimension. Then $\mathrm{Sh}(X; \mathcal{C})$ has enough points. That is, a morphism $\alpha : F \rightarrow G$ is an equivalence if and only if $\alpha_x : F_x \rightarrow G_x$ is an equivalence for all $x \in X$.*

Proof. The “only if” direction is automatic, so we show the “if” direction. We want to check that the canonical map $\alpha(U) : F(U) \rightarrow G(U)$ is an equivalence for all open subsets $U \subseteq X$. Let $\{S_k\}$ be a generating family of compact objects in $\mathcal{C}$. By the previous lemma, this holds if and only if

$$\alpha(U) \circ (-) : \mathrm{Hom}_{\mathcal{C}}(S_k, F(U)) \rightarrow \mathrm{Hom}_{\mathcal{C}}(S_k, G(U))$$

is an equivalence for all S_k and all U . Now, if $F \in \mathrm{Sh}(X; \mathcal{C})$, then composing with $\mathrm{Hom}(C, -)$ gives

$$\begin{aligned} F_C : \mathrm{Op}_X^{\mathrm{op}} &\rightarrow \mathrm{Ani} \\ U &\mapsto \mathrm{Hom}(C, F(U)), \end{aligned}$$

which defines an anima-valued sheaf for any object $C \in \mathcal{C}$. If we furthermore assume C to be compact, then

$$(F_C)_x = \mathrm{colim}_{U \ni x} \mathrm{Hom}(C, F(U)) = \mathrm{Hom}(C, \mathrm{colim}_{U \ni x} F(U)) = \mathrm{Hom}(C, F_x),$$

where the second equality uses that C is compact. Thus, $\alpha_x : F_x \rightarrow G_x$ is an equivalence if and only if $\alpha_x : (F_{S_k})_x \rightarrow (G_{S_k})_x$ is an equivalence for all k . Applying the case of anima-valued sheaves (Theorem 3.1.2), we conclude that $F_{S_k} \rightarrow G_{S_k}$ is an equivalence for all k , which is what we wanted to show. $\square$

Having the ability to check at stalks implies a special case of proper base change.

Lemma 3.1.14. ²⁶ *Let $\mathcal{C}$ be an ω -accessible presentable category and let X be of finite dimension. Let $i : Z \subseteq X$ be a closed subset and let $f : Y \rightarrow X$ be a continuous map. Consider the pullback diagram*

$$\begin{array}{ccc} f^{-1}(Z) & \xrightarrow{f'} & Z \\ \downarrow i' & & \downarrow i \\ Y & \xrightarrow{f} & X \end{array}$$

*in topological spaces. Then the canonical map $f^*i_* \rightarrow i'_*f'^*$ is an equivalence. In particular, taking $Y \rightarrow X$ to be the inclusion of a point $\{x\} \hookrightarrow X$, we have, for any $F \in \mathrm{Sh}(Z; \mathcal{C})$,*

$$(i_*F)_x = \begin{cases} F_x, & x \in Z, \\ *, & x \notin Z. \end{cases}$$

Proof. We first show that the special case implies the general case. Let $y \in Y$ and $F \in \mathrm{Sh}(Z; \mathcal{C})$. Since equivalences can be checked on stalks, we compute that if $y \in f^{-1}(Z)$,

$$(f^*i_*F)_y = (i_*F)_{f(y)} = F_{f(y)} = (i'_*f'^*F)_y.$$

The same computation shows that if $y \notin f^{-1}(Z)$, both sides are $*$. For the special case, we compute

$$(i_*F)_x = \mathrm{colim}_{U \ni x} F(U \cap Z),$$

²⁶We need to prove this lemma since we do not have proper base change outside the stable setting except for anima-valued sheaves.

where U runs over all open neighborhoods of x , and the assertion follows immediately. $\square$

Note that the stable situation is even simpler. Using, for example, the criterion that a category is stable when Ω_* is invertible (Lemma 2.4.8), we see that the equivalence restricts to stable categories.

Notation 3.1.15. We use the notation $\text{st} \subseteq \text{Cat}^{\text{Rex}}$ to denote the subcategory of small stable idempotent-complete categories with exact functors. Similarly, we denote by $\text{Pr}_{\omega, \text{st}}^{\text{L}}$ the category of stable ω -accessible presentable categories with colimit-preserving functors.

Corollary 3.1.16. *The equivalence of Proposition 3.1.9 restricts to*

$$\begin{aligned} \text{Ind} : \text{st} &\xrightarrow{\sim} \text{Pr}_{\omega, \text{st}}^{\text{L}} \\ \mathcal{C}_0 &\mapsto \text{Ind}(\mathcal{C}_0). \end{aligned}$$

Lemma 3.1.17. *Let $\mathcal{C}$ be stable. Then a family of objects $\mathcal{S} = \{S_\alpha\} \subseteq \mathcal{C}$ is a generating family if and only if $\text{Hom}_{\mathcal{C}}(S_\alpha, X) = 0$ for all α implies $X = 0$.*

Proof. Let $f : Y \rightarrow Z$ be a morphism in $\mathcal{C}$. There is a fiber sequence

$$Y \xrightarrow{f} Z \rightarrow \text{cof}(f).$$

The implication “ $\Leftarrow$ ” follows from the fact that f is an equivalence if and only if $\text{cof}(f) = 0$. For the implication “ $\Rightarrow$ ”, one uses the fact that any object admits a canonical map $X \rightarrow 0$. $\square$

Corollary 3.1.18. *Let $\mathcal{C}$ be a stable ω -accessible presentable category and let X be of finite dimension. Then a sheaf $F \in \text{Sh}(X; \mathcal{C})$ is zero if and only if $F_x = 0$ for all $x \in X$.*

We now turn our attention to the notion of support. Let $\mathcal{C}$ be a category with a final object, let $j : U \hookrightarrow X$ be an open subset, and let $F \in \text{Sh}(X; \mathcal{C})$ be a sheaf. We use the notation $F|_U := j^*F$ for its restriction. Denote by $*_X$ the sheaf that sends every open set of X to the final object $*$. The universal property of $*$ in $\mathcal{C}$ then implies that, for any $F \in \text{Sh}(X; \mathcal{C})$, there exists a unique map $F \rightarrow *_X$, so $*_X$ is the final object in $\text{Sh}(X; \mathcal{C})$.

Lemma 3.1.19. *Let $\mathcal{C}$ be a category with a final object $*$. For $F \in \text{Sh}(X; \mathcal{C})$, if U is an open set and $\{U_\alpha\}$ is an open cover of U such that the unique map $F|_{U_\alpha} \rightarrow *_{U_\alpha}$ is an equivalence for each α , then the unique map $F|_U \rightarrow *_U$ is an equivalence. In particular, there is a largest open set U such that $F|_U \rightarrow *_U$ is an equivalence.*

Proof. This is a direct consequence of the sheaf condition (1). $\square$

Definition 3.1.20. Let $\mathcal{C}$ be a category with a final object, let X be a topological space, and let $F \in \text{Sh}(X; \mathcal{C})$ be a sheaf. The support of F , denoted $\text{supp}(F)$, is the closed complement of the largest open set $U \subseteq X$ such that $F|_U = *_U$.

Remark 3.1.21. For those familiar with the abelian or stable setting, $\text{supp}(F)$ should be thought of as the locus where the sheaf is “non-trivial.” A natural question is why triviality should be defined via the final object $*$ rather than, say, the initial object $\emptyset$, when $\mathcal{C}$ admits both and they differ. One consideration is the following. The empty set $j : \emptyset \subseteq X$ is an open set and $\text{Sh}(\emptyset; \mathcal{C}) = \{*_\emptyset\}$: since $\text{Op}_\emptyset = \{\emptyset\}$ contains exactly one object, the sheaf condition (Remark 2.3.4) forces any sheaf to take the value $*$ on $\emptyset$. One then computes directly that $j_*(*_\emptyset) = *_X$, so the sheaf coming from the empty set is $*_X$.

Lemma 3.1.22. *Let $Z \subseteq X$ be a closed subset. The full subcategory*

$$\{F \in \mathrm{Sh}(X; \mathcal{C}) \mid \mathrm{supp}(F) \subseteq Z\} \subseteq \mathrm{Sh}(X; \mathcal{C})$$

is closed under limits that exist in $\mathrm{Sh}(X; \mathcal{C})$. If $\mathcal{C}$ is pointed, then it is also closed under colimits that exist in $\mathrm{Sh}(X; \mathcal{C})$.

Proof. This is a direct consequence of unwrapping the definitions and noting that, for any open embedding $j : U \hookrightarrow X$, the functor j^* preserves all existing limits and colimits. $\square$

Proposition 3.1.23. *Let $\mathcal{C}$ be a category with a final object and let $Z \subseteq X$ be a closed subset. Then $i_* : \mathrm{Sh}(Z; \mathcal{C}) \hookrightarrow \mathrm{Sh}(X; \mathcal{C})$ is fully faithful. If we further assume that $\mathcal{C}$ is an ω -accessible presentable category and X is finite-dimensional, then the image of i_* yields an equivalence*

$$i_* : \mathrm{Sh}(Z; \mathcal{C}) \xrightarrow{\sim} \{F \in \mathrm{Sh}(X; \mathcal{C}) \mid \mathrm{supp}(F) \subseteq Z\}.$$

Proof. For a functor $f : \mathcal{C} \rightarrow \mathcal{D}$ admitting a left adjoint, being fully faithful is equivalent to the counit $f^L f \rightarrow \mathrm{id}_{\mathcal{C}}$ being an equivalence. Full faithfulness in our case therefore follows from Example 2.3.16. To see the characterization of the image, set $U := Z^c$ and denote by $j : U \hookrightarrow X$ the inclusion. Lemma 3.1.14 then implies $j^* i_* F = *_{\mathcal{U}}$, so $U \subseteq \mathrm{supp}(i_* F)^c$ and hence $\mathrm{supp}(i_* F) \subseteq Z$. For the reverse inclusion, let $G \in \mathrm{Sh}(X; \mathcal{C})$ with $\mathrm{supp}(G) \subseteq Z$ and consider the unit map $G \rightarrow i_* i^* G$. One can then check at stalks and use the special case of Lemma 3.1.14 to conclude that it is an equivalence. $\square$

Lemma 3.1.24. *Let $\mathcal{C}$ be a category with a final object. Then $\overline{\{x \mid F_x \neq *\}} \subseteq \mathrm{supp}(F)$. If we further assume that $\mathcal{C}$ is an ω -accessible presentable category and X is finite-dimensional, then $\mathrm{supp}(F) = \overline{\{x \mid F_x \neq *\}}$.*

Proof. For the inclusion $\overline{\{x \mid F_x \neq *\}} \subseteq \mathrm{supp}(F)$: since $\mathrm{supp}(F)$ is closed, it suffices to show $\{x \mid F_x \neq *\} \subseteq \mathrm{supp}(F)$, which is equivalent to $\mathrm{supp}(F)^c \subseteq \{x \mid F_x = *\}$, and this follows directly from the definition. The reverse inclusion follows from $\mathrm{supp}(F)^c \supseteq \mathrm{Int}(\{x \mid F_x = *\})$. In other words, we need to show that if $F_x = *$ for all $x \in U$ for some open set U , then $F|_U = *_{\mathcal{U}}$, which is exactly Proposition 3.1.13. $\square$

Again, the stable situation is easier to phrase. We begin with a well-known fact.

Definition 3.1.25. Let $\mathcal{C} \in \mathrm{Pr}_{st}^L$ and $\iota_* : \mathcal{C}_0 \hookrightarrow \mathcal{C}$ and $j_! : \mathcal{C}_1 \hookrightarrow \mathcal{C}$ be subcategories. We say that $(\mathcal{C}_0, \mathcal{C}_1)$ forms a recollement of $\mathcal{C}$ if there exist adjunctions

$$\begin{array}{ccccc} & \overset{\iota^*}{\curvearrowright} & & \overset{j^!}{\curvearrowright} & \\ \mathcal{C}_0 & \xrightleftharpoons[\iota^!]{\iota_*} & \mathcal{C} & \xrightleftharpoons[j_*]{j^*} & \mathcal{C}_1 \end{array}$$

such that $j^* \iota_* = 0$ and the unit and counit of the functors form exact sequences

$$j_! j^* \rightarrow \mathrm{id}_{\mathcal{C}} \rightarrow \iota_* \iota^* \text{ and } \iota_* \iota^! \rightarrow \mathrm{id}_{\mathcal{C}} \rightarrow j_* j^*.$$

In this case, j_* is also fully faithful and one can identify $\mathcal{C}_1$ with ${}^\perp \mathcal{C}_0$, the left orthogonal complement of $\mathcal{C}_0$, via $j_!$, and with $\mathcal{C}_0^\perp$, the right orthogonal complement of $\mathcal{C}_0$, via j_* .

Remark 3.1.26. In fact, the notion of recollement exists more generally for categories with finite limits, and is also weaker. See [45, Definition A.8.1]. Nevertheless, one can show, as in [45, Remark A.8.19], that what we describe in Definition 3.1.25 is equivalent to the general notion in the stable setting.

Notation 3.1.27. Let X be a locally compact Hausdorff space and let $\mathcal{C}$ be a stable presentable category. For a locally closed subset $k : Z \hookrightarrow X$ and $F \in \mathrm{Sh}(X; \mathcal{C})$, we use the notation

$$F_Z := k_! k^* F \quad \text{and} \quad \Gamma_Z(F) := k_* k^! F.$$

Lemma 3.1.28. *Let $\mathcal{C}$ be a stable presentable category, let $i : Z \subseteq X$ be a closed subset, and let $j : U \subseteq X$ be its open complement. Then, for any $F \in \mathrm{Sh}(X; \mathcal{C})$, there is a fiber sequence*

$$F_U \rightarrow F \rightarrow F_Z,$$

where the maps are given by the counit and unit of the corresponding adjunctions. In particular, i_ induces an equivalence $\mathrm{Sh}(Z; \mathcal{C}) \xrightarrow{\sim} \{F \in \mathrm{Sh}(X; \mathcal{C}) \mid \mathrm{supp}(F) \subseteq Z\}$, and its cofiber is given by $j^* : \mathrm{Sh}(X; \mathcal{C}) \rightarrow \mathrm{Sh}(U; \mathcal{C})$. For the right adjoints, we also have the fiber sequence*

$$\Gamma_Z(F) \rightarrow F \rightarrow \Gamma_U(F).$$

In fact, the following is a rephrased version of Corollary 3.1.18.

Corollary 3.1.29. *Let $\mathcal{C}$ be a stable ω -accessible presentable category and let X be a topological space of finite dimension. Then for $F \in \mathrm{Sh}(X; \mathcal{C})$,*

$$\mathrm{supp}(F) = \overline{\{x \mid F_x \neq 0\}}.$$

Remark 3.1.30. In light of Lemma 3.1.28, it is possible that Proposition 3.1.23 is not optimal, i.e., the ω -accessible assumption may be unnecessary. However, as far as the author is aware, the existing general version in, for example, [66, Theorem 4.4] shows that the pointed setting is enough. To keep the exposition simple, and given that we will rely on stalks and similar objects quite often later, so that facts in the spirit of Lemma 3.1.24 will be unavoidable, we have chosen to take this shortcut.

3.2. Microsupport. We have considered the notion of support, which records where a sheaf is non-trivial, i.e., where $\Gamma(U; F)$ is not zero (or, more generally, final). A similar question can be asked for restriction maps: we would like invariants that detect whether the restriction map

$$\Gamma(U; F) \rightarrow \Gamma(V; F),$$

for $V \subseteq U$, is non-trivial, i.e., not an equivalence. For this purpose, we assume our space X is a manifold, so that there is a well-behaved notion of “gradual change” of open sets.

Notation 3.2.1. For a C^∞ -manifold M , we use the notation 0_M for the zero section in T^*M , and set $\dot{T}^*M := T^*M \setminus 0_M$ for the cotangent bundle away from the zero section, and $S^*M := \dot{T}^*M / \mathbb{R}^+$ for the cosphere bundle.

Definition 3.2.2. ²⁷ Let $\Omega \subseteq \dot{T}^*M$ be a conic open subset. An Ω -testing pair (U, g) consists of an open set U and a C^∞ -function $g : U \times [0, 1] \rightarrow \mathbb{R}$ such that, setting $g_t := g(-, t)$,

- (1) $d(g_t)_x \in \Omega$ for all $(x, t) \in U \times [0, 1]$,
- (2) $\{g_{t_1} < 0\} \subseteq \{g_{t_2} < 0\}$ whenever $t_1 \leq t_2$,
- (3) $\Sigma := \{g_1 < 0\} \setminus \{g_0 < 0\}$ has compact closure, and
- (4) $\{g_t = 0\} \cap (U \setminus \overline{\Sigma})$ is independent of $t \in [0, 1]$.

We call a locally closed set Σ an Ω -lens if it can be obtained as $\Sigma := \{g_1 < 0\} \setminus \{g_0 < 0\}$ for some Ω -testing pair (U, g) .

²⁷Warning: This notation is nonstandard. We choose it so that the usual definition makes sense unstably. See Remark 3.2.16.




FIGURE 1. While the restriction maps go the other way, the geometric intuition is to understand whether sections change when expanding the open sets on which they live. Here the green shaded area is $B \cap \{x < t\}$ for $t = 0, 0.5$, and 1 .

One should think of the data of U and g as a way to interpolate smoothly between the open sets $\{g_0 < 0\} \subseteq \{g_1 < 0\}$. For example, restricting g to $U \times [0, c]$ for any $c \in (0, 1)$, up to reparametrization, produces a (possibly smaller) Ω -lens. We can now give the definition of microsupport.

Definition 3.2.3. Let $F \in \text{Sh}(M; \mathcal{C})$ be a sheaf.

- (1) We say F is microlocally trivial on the conic open set Ω if, for any Ω -testing pair (U, g) , the restriction map

$$\Gamma(\{g_1 < 0\}; F) \xrightarrow{\sim} \Gamma(\{g_0 < 0\}; F)$$

is an equivalence.

- (2) We define the microsupport of F away from zero, $\dot{\text{SS}}(F) \subseteq \dot{T}^*M$, to be the complement of the union of all conic open sets $\Omega \subseteq \dot{T}^*M$ on which F is microlocally trivial. Since $\dot{\text{SS}}(F)$ is conic, it is determined by its projection

$$\text{SS}^\infty(F) := (\dot{\text{SS}}(F)/\mathbb{R}^+) \subseteq S^*M$$

to the cosphere bundle, which we call the microsupport at infinity.

- (3) We define the microsupport of F , $\text{SS}(F)$, to be the union

$$\text{SS}(F) := \text{supp}(F) \cup \dot{\text{SS}}(F) \subseteq T^*M$$

in the cotangent bundle.²⁸

Remark 3.2.4. Since we test against arbitrarily small U , we have $\pi_\infty(\text{SS}^\infty(F)) \subseteq \text{supp}(F)$, where $\pi_\infty : S^*M \rightarrow M$ is the bundle projection. Thus $\text{SS}(F)$ encodes the data of $\text{supp}(F)$ together with the co-directions of change, or “non-propagation,” of sections.

Example 3.2.5. Consider the case $M = \mathbb{R}^2$ with coordinates (x, y) and $\Omega = B \times C$, where B is the unit open ball and C is a proper open cone containing the covector $(1, 0)$, e.g., $C = \{(\xi, \eta) \mid \xi > 0, \xi < |\eta|\}$. We claim that $\Sigma := \{x \geq 0\} \cap B$ is an Ω -lens. One can check directly that $U = B$ and $g : B \times [0, 1] \rightarrow \mathbb{R}$ given by $g(x, y, t) = x - t$ give the desired open set and function. More explicitly, this pair measures the difference of the restriction map

$$\Gamma(B; F) \rightarrow \Gamma(B \cap \{x < 0\}; F).$$

See Figure 1 for an illustration.

If we further assume $\mathcal{C}$ is stable and presentable, then Lemma 3.1.28 implies that there is a fiber sequence

$$1_{\{g_0 < 0\}} \rightarrow 1_{\{g_1 < 0\}} \rightarrow 1_\Sigma.$$

²⁸The notation SS comes from the older name “singular support.”

In particular, for $F \in \text{Sh}(M; \mathcal{C})$, the map $\Gamma(\{g_1 < 0\}; F) \xrightarrow{\sim} \Gamma(\{g_0 < 0\}; F)$ is an equivalence if and only if $\text{Hom}(1_\Sigma, F) = 0$. We thus have a more concise, if psychologically awkward, equivalent definition of microlocal triviality.

Definition 3.2.6. Let $\mathcal{C}$ be a stable presentable category. We say a sheaf $F \in \text{Sh}(M; \mathcal{C})$ is microlocally trivial on the conic open set Ω if, for any Ω -lens Σ , we have $\text{Hom}(1_\Sigma, F) = 0$.

Note that $\text{Hom}(1_\Sigma, -)$ is limit-preserving, and we would like a dual condition that is colimit-preserving.

Lemma 3.2.7 ([29, Lemma 3.3]). *Let $\mathcal{C}$ be a stable presentable category, let $F \in \text{Sh}(M; \mathcal{C})$, and let $\Omega \subseteq T^*M$ be a conic open set. Then $\text{SS}(F) \cap \Omega = \emptyset$ if and only if*

$$\Gamma(M; 1_\Sigma \otimes F) = 0$$

*for every Ω^a -lens Σ , where Ω^a denotes the image of Ω under the antipodal map $T^*M \rightarrow T^*M$, $(x, \xi) \mapsto (x, -\xi)$. Moreover, by the compactness assumption on $\overline{\Sigma}$, the expression $\Gamma(M; 1_\Sigma \otimes F) = \Gamma_c(M; 1_\Sigma \otimes F)$ is colimit-preserving in F .*

Proof. Some explanation of the proof is needed, as the authors use [36, Proposition 2.5.1]²⁹ in the proof. The argument proceeds as follows. Let Σ be an Ω^a -lens and write it as $\overline{\Sigma} \setminus Z$, where $Z := \overline{\Sigma} \setminus \Sigma$. One can then find two families of open sets $\{V_n\}$ and $\{U_n\}$ such that $\bigcap U_n = \overline{\Sigma}$, $\bigcap V_n = Z$, $V_n \subseteq U_n$, and $U_n \setminus V_n$ is an Ω -lens for all n . The equivalences $\Gamma(U_n; F) \xrightarrow{\sim} \Gamma(V_n; F)$ then imply, by [36, Proposition 2.5.1], that $\Gamma(\overline{\Sigma}; F|_{\overline{\Sigma}}) \xrightarrow{\sim} \Gamma(Z; F|_Z)$. Now, for any closed subset $X \subseteq M$, one has $\Gamma(X; F|_X) = \Gamma(M; F \otimes 1_X)$, so taking the fiber yields $\Gamma(M; 1_\Sigma \otimes F) = 0$.

The use of [36, Proposition 2.5.1] can be avoided as follows. Since we are working on a manifold, by carefully shrinking U_n and V_n inductively, one can further arrange that $U_{n+1} \subseteq \overline{U_{n+1}} \subseteq U_n$ and $\overline{U_n}$ is compact, and similarly for the V_n 's. The transition map $\Gamma(U_n; F) \rightarrow \Gamma(U_{n+1}; F)$ then factors as

$$\Gamma(U_n; F) \rightarrow \Gamma(\overline{U_{n+1}}; F) \rightarrow \Gamma(U_{n+1}; F).$$

Thus we have

$$\begin{aligned} \text{colim}_{n \rightarrow \infty} \Gamma(U_n; F) &= \text{colim}_{n \rightarrow \infty} \Gamma(\overline{U_n}; F) = \text{colim}_{n \rightarrow \infty} \Gamma(M; F \otimes 1_{\overline{U_n}}) \\ &= \Gamma\left(M; F \otimes \left(\text{colim}_{n \rightarrow \infty} 1_{\overline{U_n}}\right)\right) = \Gamma(M; F \otimes 1_{\overline{\Sigma}}), \end{aligned}$$

as desired. Here we use the fact that $\text{supp}(F \otimes 1_{\overline{U_n}}) \subseteq \overline{U_1}$ is compact to pass the colimit inside in the penultimate equality. $\square$

Remark 3.2.8. We claim that if $\mathcal{C}$ is a stable ω -accessible presentable category, then the following holds: for X locally compact Hausdorff, $K \subseteq X$ compact, and $F \in \text{Sh}(X; \mathcal{C})$, the canonical map

$$\text{colim}_{U \supseteq K} \Gamma(U; F) \rightarrow \Gamma(K; F)$$

is an equivalence. The case $\mathcal{C} = \text{Ab}$ is [36, Proposition 2.5.1]. As in the proof of Proposition 3.1.13, by testing against compact generators one may assume $\mathcal{C} = \text{Sp}$. By part (3) of [45, Proposition 1.4.3.6], there is a standard t -structure on Sp , and the standard argument reduces the claim to the case $\mathcal{C} = \text{Ab}$.

²⁹We give the explanation of why it holds for stable ω -accessible presentable categories in Remark 3.2.8 below.

Notation 3.2.9. For a conic closed set $X \subseteq T^*M$, we use the notation

$$\mathrm{Sh}_X(M) := \{F \mid \mathrm{SS}(F) \subseteq X\}$$

to denote the subcategory of sheaves microsupported in X . For a closed set $\bar{X} \subseteq S^*M$, we set

$$\mathrm{Sh}_{\bar{X}}(M) := \mathrm{Sh}_{\mathbb{R}^+ \bar{X} \cup 0_M}(M),$$

where $\mathbb{R}^+ \bar{X} \subseteq T^*M$ is the unique conic subset projecting to $\bar{X}$ in S^*M .

Example 3.2.10. We have an equality $\mathrm{Loc}(M) = \mathrm{Sh}_{0_M}(M)$. For example, locally on a chart with coordinates $(x_1, \dots, x_n)$, one can take $g_t = (x_1^2 + \dots + x_n^2) + \epsilon - t$, which implies that $\Gamma(B_{t+\epsilon}(x); F) \xrightarrow{\sim} \Gamma(B_\epsilon(x); F)$ is an equivalence for arbitrarily small $\epsilon > 0$.

Example 3.2.11. Let $Z \subseteq M$ be a closed subset and set $T^*M|_Z := \{(x, \xi) \in T^*M \mid x \in Z\}$. Then $\mathrm{Sh}_{T^*M|_Z}(M) = \{F \mid \mathrm{supp}(F) \subseteq Z\}$.

Let $Z \subseteq M$ be a closed subset. By Lemma 2.5.11, since we assume $\mathcal{C}$ to be presentable, the category $\mathrm{Sh}(M; \mathcal{C})$ is presentable. As a result, the category $\{F \mid \mathrm{supp}(F) \subseteq Z\}$ is presentable by Proposition 3.1.23, since it is equivalent to $\mathrm{Sh}(Z; \mathcal{C})$. For a conic closed set $X \subseteq T^*M$, there is no simple “classical” expression of this form. Nevertheless, Definition 3.2.6 and Lemma 3.2.7 imply that $\mathrm{Sh}_X(M; \mathcal{C}) \subseteq \mathrm{Sh}(M; \mathcal{C})$ is closed under both limits and colimits, so presentability follows from the theorem below:

Theorem 3.2.12 ([53]). *Let $\mathcal{C}$ be a presentable category. If an inclusion $\mathcal{D} \subseteq \mathcal{C}$ is limit-preserving and accessible, i.e., preserving κ -filtered colimits for some regular cardinal κ , then $\mathcal{D}$ is presentable.*

Corollary 3.2.13. *Let $X \subseteq T^*M$ be a conic closed subset. Then the category $\mathrm{Sh}_X(M)$ of sheaves microsupported in X is presentable. As a result, the inclusion $\iota_X : \mathrm{Sh}_X(M) \hookrightarrow \mathrm{Sh}(M)$ admits both a left adjoint and a right adjoint, which we denote by ι_X^* and $\iota_X^!$ respectively.*

Similarly to the situation of support in Lemma 3.1.24 and Corollary 3.1.29, there are both open-set and stalk viewpoints.

Definition 3.2.14 ([26, Section 4.3]). Let $F \in \mathrm{Sh}(M; \mathcal{C})$ be a sheaf and let ϕ be a C^∞ -function defined on some open set $U \subseteq M$. Denote by $i_{\phi, t} : \{\phi(x) \leq t\} \subseteq U$ the closed inclusion. We say $x \in M$ is a cohomological F -critical point of ϕ at level t if $(i_{\phi, t}^! F)_x \neq 0$.

Proposition 3.2.15 ([29, Lemma 3.2]). *Assume $\mathcal{C}$ is a stable ω -accessible presentable category. Then, for $F \in \mathrm{Sh}(M; \mathcal{C})$, the microsupport of F is the closure of the locus of differentials of C^∞ -functions at their cohomological F -critical points. That is,*

$$\mathrm{SS}(F) = \overline{\bigcup_{\phi \in C^\infty} \{(x, \xi) \mid \exists t \in \mathbb{R}, (i_{\phi, t}^! F)_x \neq 0, \xi = d\phi_x\}}.$$

Remark 3.2.16. The most standard definition of microsupport, given by Kashiwara and Schapira as [36, (1) of Proposition 5.1.1], is the following. Let $F \in \mathrm{Sh}(M; \mathcal{C})$ be a sheaf. The microsupport $\mathrm{SS}(F)$ is the complement of those $(x, \xi) \in T^*M$ such that the following holds: there exists an open set $\Omega \ni (x, \xi)$ such that, for any $y \in M$ and any C^∞ -function ϕ defined near y satisfying $\phi(y) = 0$ and $d\phi_y \in \Omega$, we have $(i_{\phi, 0}^! F)_y = 0$. The equivalent description in Proposition 3.2.15, taken from [26, Section 4.3], is the same definition stated directly from

the viewpoint of $\mathrm{SS}(F)$ rather than its complement. It is proved in [29, Lemma 3.2] that the definition given here is equivalent to the standard one when $\mathcal{C}$ is ω -accessible presentable.

There are two main reasons for this choice of presentation. First, the definition as presented is what makes Corollary 3.2.13 hold. From the technical viewpoint, in the higher categorical setting it is very difficult to work with a category whose limit- or colimit-preservation properties are unknown. The more fundamental reason, from the author's perspective, is the failure of the noncharacteristic deformation lemma ([36, Proposition 2.7.2]) when the coefficient $\mathcal{C}$ is presentable but not ω -accessible, as observed by Efimov in [18, Remark 4.23]. Roughly speaking, in the simplest case, it says that if $\{U_t\}_{t \in [0,1]}$ is a family of expanding open sets with smooth boundary that changes in such a way that the outward conormal $N_{\mathrm{out}}^*(U_t)$ avoids $\mathrm{SS}(F)$ as t increases, then the restriction map

$$\Gamma(U_1; F) \xrightarrow{\sim} \Gamma(U_0; F)$$

should be an equivalence.

In light of Remark 3.2.16, we will henceforth make the following assumption. For the purposes of applications, ω -accessibility is, as far as the author is aware, almost always available. Moreover, compatibility between microsupport and the six-functor formalism, which is the next topic of discussion, is only available in the literature under the classical definition.

Assumption 3.2.17. Unless otherwise specified, we fix a stable ω -accessible presentable category $\mathcal{V}$. As a result, we adopt the notation $\mathrm{Sh}(M) := \mathrm{Sh}(M; \mathcal{V})$, as anima-valued sheaves $\mathrm{Sh}(M; \mathrm{Ani})$ will seldom be discussed further. Similarly, all functors should be $\mathcal{V}$ -linear, so instead of writing $\mathrm{Fun}_{\mathcal{V}}^L$ or $\otimes_{\mathcal{V}}$, we will use the simplified notations Fun^L and $\otimes$.

Regarding the underlying geometry, and for the purposes of microsupport, all maps between manifolds $f : M \rightarrow N$ are assumed to be C^∞ .

3.3. Compatibility and simple examples. We recollect, without proof, the compatibility of microsupport with various natural operations on sheaves. One often begins with a sheaf of simple microsupport and tries to track the possible microsupport that may result from various constructions. We begin with a rephrasing of the limit- and colimit-preserving part of Corollary 3.2.13.

Lemma 3.3.1. *Let $F : I \rightarrow \mathrm{Sh}(M; \mathcal{C})$ be a diagram of sheaves. Then we have the following inclusions:*

$$\mathrm{SS}\left(\mathrm{colim}_{\alpha \in I} F_\alpha\right) \subseteq \overline{\bigcup_{\alpha \in I} \mathrm{SS}(F_\alpha)} \quad \text{and} \quad \mathrm{SS}\left(\lim_{\alpha \in I} F_\alpha\right) \subseteq \overline{\bigcup_{\alpha \in I} \mathrm{SS}(F_\alpha)}.$$

Proof. This follows from Corollary 3.2.13 by taking $X = \overline{\bigcup_{\alpha \in I} \mathrm{SS}(F_\alpha)}$ and noting that $\mathrm{SS}(F_\alpha) \subseteq X$ for all $\alpha \in I$ tautologically. For the converse, if $\mathrm{SS}(F_\alpha) \subseteq X$ for some conic closed set $X \subseteq T^*M$, then

$$\mathrm{SS}\left(\mathrm{colim}_{\alpha \in I} F_\alpha\right) \subseteq \overline{\bigcup_{\alpha \in I} \mathrm{SS}(F_\alpha)} \subseteq \overline{\bigcup_{\alpha \in I} X} = X.$$

□

A special case of the above is that if $F' \rightarrow F \rightarrow F''$ is a fiber sequence, then $\mathrm{SS}(F'') \subseteq \mathrm{SS}(F) \cup \mathrm{SS}(F')$, since $F'' = \mathrm{cof}(F' \rightarrow F)$ is a colimit. The same bound holds for F , since




FIGURE 2. The green region is the closed cone in $\mathbb{R}^2$ given by the sector between angles 0 and $3\pi/4$. Since the statement is pointwise and, at the boundary, the region is locally modeled by a closed half-plane, which is a closed convex cone whose dual cone is the inward-pointing perpendicular ray, we draw red hairs along the two edges to indicate the microsupport there. One observes that these hairs converge at the origin, where the dual cone, represented by the fan of longer hairs, spans the range of angles from $\pi/4$ to $\pi/2$.

$F' \rightarrow F \rightarrow F''$ is a fiber sequence if and only if $F''[-1] \rightarrow F' \rightarrow F$ is a fiber sequence. There is also a lower bound.

Proposition 3.3.2 ([36, Proposition 5.1.3], the triangular inequalities). *Let $F' \rightarrow F \rightarrow F''$ be a fiber sequence. Then*

$$(\mathrm{SS}(F'') \setminus \mathrm{SS}(F')) \cup (\mathrm{SS}(F') \setminus \mathrm{SS}(F'')) \subseteq \mathrm{SS}(F) \subseteq \mathrm{SS}(F') \cup \mathrm{SS}(F'').$$

This is commonly referred to as the microlocal triangular inequalities.

Example 3.3.3 ([36, Proposition 5.3.1]). Let V be a finite-dimensional Euclidean vector space, let $\gamma \subseteq V$ be a closed convex cone with vertex at 0, and denote by $\gamma^\circ := \{\xi \in V^\vee \mid \xi(v) \geq 0 \ \forall v \in \gamma\}$ its dual cone. Then we have

$$\mathrm{SS}(1_\gamma) \cap T_0^*V = \gamma^\circ.$$

Note that a linear subspace $L \subseteq V$ is such a case, and this implies that if $M' \subseteq M$ is a closed submanifold, then $\mathrm{SS}(1_{M'}) = N^*(M')$ is the conormal bundle of M' .

Similarly, if $U \subseteq M$ is an open set with smooth boundary, i.e., ∂U is a closed submanifold of codimension 1, then $\overline{U}$ near $x \in \partial U$ is modeled by a closed half-space $\{v \in V \mid \xi(v) \leq 0\}$ for some $\xi \neq 0$ in $V^\vee$, which is a closed convex cone. Thus we have $\mathrm{SS}(1_{\overline{U}}) = N_{in}^*(\overline{U}) := \overline{U} \cup N_{in}^*(\partial U)$, where the inward conormal $N_{in}^*(\partial U) \subseteq N^*(\partial U)$ consists of those ξ pairing non-negatively with vectors pointing into U . Applying the microlocal triangular inequalities (Proposition 3.3.2) to the fiber sequence $1_U \rightarrow 1_M \rightarrow 1_{M \setminus U}$, we obtain $\mathrm{SS}(1_U) = N_{out}^*(U)$.

In order to state the compatibility between microsupport and pushforward and pullback, we introduce some notation.

Notation 3.3.4. Let $f : M \rightarrow N$ be a C^∞ map. Then there is an induced diagram:

$$\begin{array}{ccccc} T^*M & \xleftarrow{df^\vee} & f^*T^*N & \xrightarrow{f_\pi} & T^*N \\ & & \downarrow & & \downarrow \\ & & M & \xrightarrow{f} & N \end{array}$$

Here, f^*T^*N is the pullback bundle, which can be concretely described as the set-theoretic pullback $f^*T^*N = \{(x, \eta) \mid x \in M, \eta \in T_{f(x)}^*N\}$, and the maps are given by

$$T^*M \xleftarrow{df^\vee} f^*T^*N \xrightarrow{f_\pi} T^*N$$

$$(x, \eta \circ df_x) \longleftarrow (x, \eta) \longrightarrow (f(x), \eta)$$

Example 3.3.5. Note that any C^∞ map $f : M \rightarrow N$ can be factored as an inclusion of a closed submanifold followed by a smooth projection:

$$M \xrightarrow{\Gamma_f} M \times N \rightarrow N.$$

Thus, to understand the geometry of the above notation, it usually suffices to understand the case of the standard inclusion and projection $i : \mathbb{R}^m \hookrightarrow \mathbb{R}^{m+k}$ and $p : \mathbb{R}^{m+k} \rightarrow \mathbb{R}^k$.

Choose coordinates for $T^*\mathbb{R}^{m+k}$ given by (x, y, ξ, η) . Then $i^*T^*\mathbb{R}^{m+k} \cong \mathbb{R}^m \times (\mathbb{R}^{m+k})^\vee$ has coordinates (x, ξ, η) , and $p^*T^*\mathbb{R}^k \cong \mathbb{R}^{m+k} \times (\mathbb{R}^k)^\vee$ has coordinates (x, y, η) . Furthermore, the maps are given by

$$T^*\mathbb{R}^m \xleftarrow{di^\vee} T^*\mathbb{R}^{m+k}|_{\mathbb{R}^m} \xleftarrow{i_\pi} T^*\mathbb{R}^{m+k}$$

$$(x, \xi) \longleftarrow (x, \xi, \eta) \longrightarrow (x, 0, \xi, \eta)$$

and

$$T^*\mathbb{R}^{m+k} \xleftarrow{dp^\vee} \mathbb{R}^m \times T^*\mathbb{R}^k \xrightarrow{p_\pi} T^*\mathbb{R}^k$$

$$(x, y, 0, \eta) \longleftarrow (x, y, \eta) \longrightarrow (y, \eta).$$

We first consider pullbacks.

Notation 3.3.6. Let $\mathcal{C}$ be a category with small limits and colimits. For a continuous map $f : X \rightarrow Y$ in LCH, we set $\omega_f := f^!1_Y$. In the case $Y = *$, we simply write $\omega_X := \omega_f$.

Lemma 3.3.7 ([36, Proposition 3.3.2]). *Let $\mathcal{C}$ be a category with small limits and colimits. For a continuous map $f : X \rightarrow Y$ in LCH, there is a canonical map*

$$(12) \quad f^*F \otimes \omega_f \rightarrow f^!F.$$

When f is a topological submersion, i.e., f is locally of the form $Y \times \mathbb{R}^n \rightarrow Y$, the canonical map (12) is an equivalence and ω_f is an invertible local system. In particular, if M is a manifold, then ω_M is an invertible local system, locally of the form $1_U[\dim M]$ for small opens $U \subseteq M$.

Proof. We compute that

$$\mathrm{Hom}(f^*F \otimes \omega_f, f^!F) = \mathrm{Hom}(f_!(f^*F \otimes \omega_f), F) = \mathrm{Hom}(F \otimes (f_!f^!1_Y), F).$$

There is a canonical map $F \otimes (f_!f^!1_Y) \rightarrow F$ given by tensoring the counit $f_!f^!1_Y \rightarrow 1_Y$ with F . To see that (12) is an equivalence, one first shows, by induction on n , that $\omega_{\mathbb{R}^n} = 1_{\mathbb{R}^n}[n]$. For the general case, consider the projection $p_i : X_1 \times X_2 \rightarrow X_i$. There is a base change map $p_2^*\omega_{X_2} \rightarrow p_1^!1_{X_1}$, which is not an equivalence in general, but one checks on a basis that it is when $X_2 = \mathbb{R}^n$. $\square$



FIGURE 3. Side-by-side depiction of the microsupports $\text{SS}(F)$ and $\text{SS}(G)$: the line $y = x$ and the parabola $y = x^2$ (in green) together with their conormal bundles (in red).

Proposition 3.3.8 ([36, Proposition 5.4.5]). *Let $f : M \rightarrow N$ be smooth and let $F \in \text{Sh}(N)$. Then we have the equality*

$$\text{SS}(f^*F) = \text{SS}(f^!F) = df^\vee(f_\pi^{-1}(\text{SS}(F))).$$

*In fact, $G \in \text{Sh}(M)$ satisfies $\text{SS}(G) \subseteq df^\vee(f^*T^*N)$ if and only if locally on M there exists F such that $G = f^*F$.*

To treat the general case, we need a notion of transversality.

Definition 3.3.9 ([36, Definition 5.4.12]). Let $f : M \rightarrow N$ be a smooth map. We define the conormal bundle of f by the short exact sequence

$$0 \rightarrow T_M^*N \rightarrow f^*T^*N \xrightarrow{df^\vee} T^*M \rightarrow 0$$

of bundles over M . Concretely, T_M^*N consists of those (x, η) such that $\eta \circ df_x = 0$. Let A be a closed conic subset of T^*N . We say f is *noncharacteristic* for A if

$$f_\pi^{-1}(A) \cap T_M^*N \subseteq M \times_N 0_N.$$

For a sheaf $F \in \text{Sh}(N)$, we say f is noncharacteristic for F if it is noncharacteristic for $\text{SS}(F)$.

Example 3.3.10. Consider the inclusion $i : \mathbb{R}^m \hookrightarrow \mathbb{R}^{m+k}$ from Example 3.3.5. Then $T_{\mathbb{R}^m}^*\mathbb{R}^{m+k} = \{0\}^m \times (\mathbb{R}^k)^\vee$. Thus, for an inclusion of a submanifold $i : M' \hookrightarrow M$, the conormal bundle $T_{M'}^*M = N^*(M')$ is the usual conormal bundle.

Proposition 3.3.11 ([36, Proposition 5.4.13]). *Let $F \in \text{Sh}(N)$ and let $f : M \rightarrow N$ be a map noncharacteristic for F . Then:*

- (1) $\text{SS}(f^*F) \subseteq df^\vee(f_\pi^{-1}(\text{SS}(F)))$,
- (2) *the canonical map $f^*F \otimes_{\omega_f} f^!F$ is an equivalence.*

Example 3.3.12. Identify $\mathbb{R}^1$ with the x -axis in $\mathbb{R}^2$ via the inclusion $i : \mathbb{R}^1 \hookrightarrow \mathbb{R}^2$. A subset $A \subseteq T^*\mathbb{R}^2$ is noncharacteristic for $\mathbb{R}^1$ if it does not intersect the conormal of $\mathbb{R}^1$. In particular, if $A = N^*(C)$ is the conormal of a curve C , then this holds when C is transverse to $\mathbb{R}^1$. We consider the sheaves $F = 1_{\{y=x\}}$ and $G = 1_{\{y=x^2\}}$, where the former is noncharacteristic for $\mathbb{R}^1$ and the latter is not. See Figure 3.

Both sheaves restrict to $i^*F = i^*G = 1_{\{0\}} \in \text{Sh}(\mathbb{R}^1)$, the skyscraper sheaf at 0, which has microsupport $N^*(\{0\}) = \{(0, \xi) \mid \xi \in \mathbb{R}\}$. We use Example 3.3.5 to compute what the




FIGURE 4. Following the recipe, to estimate the microsupport of a pullback one first restricts to the submanifold — in this case, intersecting with the x -axis (in black) — looks at the microsupport (in red) there, and then projects out the conormal direction. For example, in the case of F , the relevant microsupport over the intersection point is the line $\{(\xi, -\xi) \mid \xi \in \mathbb{R}\}$, which projects to the entire ξ -axis on the cotangent fiber at 0.

microsupport estimate of Proposition 3.3.11, namely $df^\vee(f_\pi^{-1}(\text{SS}(F)))$, predicts, and to verify that it fails for G . Restricting to $\mathbb{R}^1$, the conormal of $\{y = x\}$ restricts to $\{(\xi, -\xi) \mid \xi \in \mathbb{R}\}$ and the conormal of $\{y = x^2\}$ restricts to $\{(0, \xi) \mid \xi \in \mathbb{R}\}$. Projecting along $(\xi, \eta) \mapsto \xi$ then yields $N^*(\{0\})$ in the case of F and $\{0\}$ in the case of G , confirming that the noncharacteristic assumption is necessary. See Figure 4 for an illustration.

We mention that, as dual to the situation of smooth maps in Lemma 3.3.7, we have the following invertibility result for ω_f when f is the embedding of a closed submanifold.

Lemma 3.3.13. *Let $i : Z \hookrightarrow M$ be an embedding of a closed submanifold of codimension l . The sheaf $\omega_i := i^!1_M$ is a local system which is locally of the form $1_Z[-l]$.*

Proof. We may assume that the inclusion is of the form $i : \mathbb{R}^k \hookrightarrow \mathbb{R}^k \times \mathbb{R}^l$ as the first factor. On the other hand, if we let $p : \mathbb{R}^k \times \mathbb{R}^l \rightarrow \mathbb{R}^k$ be the projection, we will have $(p \circ i) = \text{id}_{\mathbb{R}^k}$, so $i^!p^!1_{\mathbb{R}^k} = 1_{\mathbb{R}^k}$. But, by Lemma 3.3.7, we have $p^!1_{\mathbb{R}^k} = 1_{\mathbb{R}^k \times \mathbb{R}^l}[l]$ and hence the result. $\square$

We consider also the case of pushforward. This time, the needed assumption is properness.

Proposition 3.3.14 ([36, Proposition 5.4.4]). *Let $f : M \rightarrow N$ be a map and let $F \in \text{Sh}(M)$ be such that f is proper on $\text{supp}(F)$. Then*

$$\text{SS}(f_*F) \subseteq f_\pi((df^\vee)^{-1}(\text{SS}(F))).$$

When f is the embedding of a closed submanifold, equality holds.

Example 3.3.15. Consider the projection $p : \mathbb{R}^2 \rightarrow \mathbb{R}^1$ given by $p(x, y) = x$ and the constant sheaf $F = 1_D$, where $D := \{x^2 + y^2 \leq 1\}$ is the closed unit disk. By Example 3.3.3, the microsupport $\text{SS}(F)$ is given by the inward conormal $N_{in}^*(D)$. Since the formula in Proposition 3.3.14, namely $f_\pi((df^\vee)^{-1}(\text{SS}(F)))$, has the reversed order compared to Proposition 3.3.11, we use Example 3.3.5 and observe that one first intersects in the cotangent direction and then projects in the base direction. This yields the microsupport estimate $N^*([-1, 1])$, which coincides with $\text{SS}(p_*F)$, since $p_*F = 1_{[-1, 1]}$. See Figure 5 for an illustration.




FIGURE 5. Following the recipe, to estimate the microsupport for pushforward one first intersects in the cotangent direction (in purple), and then projects onto the base.

We will see in Example 3.3.17 below that properness is essential to the microsupport estimate, and a remedy is stated in the following corollary.

Corollary 3.3.16. *Let M be a manifold, let $f : M \rightarrow N$ be a map between manifolds, and let $\{U_n\}$ be an increasing family of open sets such that $\bigcup_{n \in \mathbb{N}} U_n = M$ and $f|_{\overline{U_n}}$ is proper for all n . Then we have the microsupport estimates*

$$\mathrm{SS}(p_! F) \subseteq \bigcap_{n \in \mathbb{N}} \overline{\bigcup_{k \geq n} f_\pi((df^\vee)^{-1}(\mathrm{SS}(F_{U_k})))}$$

and

$$\mathrm{SS}(p_* F) \subseteq \bigcap_{n \in \mathbb{N}} \overline{\bigcup_{k \geq n} f_\pi((df^\vee)^{-1}(\mathrm{SS}(\Gamma_{U_k}(F))))}.$$

(See Notation 3.1.27 for the notation.)

Proof. Since $\mathrm{colim}_{n \rightarrow \infty} F_{U_n} \xrightarrow{\sim} F$ and $F \xrightarrow{\sim} \lim_n \Gamma_{U_n}(F)$, this is a direct consequence of Lemma 3.3.1 and Proposition 3.3.14. $\square$

Example 3.3.17. Consider again the projection $p : \mathbb{R}^2 \rightarrow \mathbb{R}^1$. Set $S := \{(x, y) \mid x > 0, y > \frac{1}{x}\}$, the open region above the graph of $y = 1/x$ in the first quadrant, and let $F = 1_S$. Base change and the identity $\Gamma_c(\mathbb{R}^n; 1_{\mathbb{R}^n}) = 1[-n]$, which follows from Poincaré duality, imply that $p_! F = 1_{(0, \infty)}[-1]$, whose microsupport $\mathrm{SS}(p_! F) = N_{\mathrm{out}}^*(0, \infty)$ contains the covector $(0, -1) \in T^*\mathbb{R}^1$. However, the conormal of the curve $y = 1/x$ at $(x, \frac{1}{x})$ is given by the covector $(\frac{-1}{x^2}, -1)$, so there are no horizontal covectors in $f_\pi((df^\vee)^{-1}(\mathrm{SS}(F)))$ (though they approach one as $x \rightarrow \infty$), and the naive formula does not detect the covector $(0, -1)$.

The remedy is to cut the support of F to F_{U_n} using $U_n := S \cap \{y < n\}$ and apply Corollary 3.3.16. More concretely, by Example 3.3.3, there is a closed cone of microsupport at $(\frac{1}{n}, n)$ interpolating between the directions $(0, 1)$ and $(-n^2, -1)$, which in particular contains $(-1, 0) \in T_{(\frac{1}{n}, n)}^*\mathbb{R}^2$. The microsupport estimate then predicts that the covector $(0, -1) \in T^*\mathbb{R}$, lying in the closure of $\{(\frac{1}{n}, -1)\}$, is a possible covector in the microsupport, which is indeed the case. See Figure 6 for an illustration.

Proposition 3.3.18 ([36, Proposition 5.4.14]). *Let $F, G \in \mathrm{Sh}(M)$ and denote by $(-)^a : T^*M \rightarrow T^*M$ the antipodal map $(x, \xi) \mapsto (x, -\xi)$.*

(1) *If $\mathrm{SS}(F) \cap (\mathrm{SS}(G))^a \subseteq 0_M$, then*

$$\mathrm{SS}(G \otimes F) \subseteq \mathrm{SS}(F) + \mathrm{SS}(G).$$




FIGURE 6. While the original region (on the left) has no horizontal outward conormal vector, after cutting it to the part inside $\{y < n\}$, a cone appears at the upper left corner, allowing the detection of the relevant covectors in the microsupport of the $!$ -pushforward.

(2) If $\text{SS}(F) \cap \text{SS}(G) \subseteq 0_M$, then

$$\text{SS}(\mathcal{H}\text{om}(G, F)) \subseteq \text{SS}(F) + (\text{SS}(G))^a.$$

Here, $+$ denotes the fiberwise sum.

3.4. Actions by isotopies. We now explicitly enter the symplectic/contact aspect of the theory. Given a manifold M , the goal of this section is to explain how a homogeneous symplectic isotopy on $\dot{T}^*M$ acts on $\text{Sh}(M)$, where $\dot{T}^*M := T^*M \setminus 0_M$ is the part of the cotangent bundle away from the zero section. Such an action is obtained via the correspondence between colimit-preserving functors and kernels: Recall that if $a \in (\mathcal{C}, \otimes)$ is dualizable with a dual $a^\vee$, then the internal Hom object $\underline{\text{Hom}}(a, b)$ exists for all $b \in \mathcal{C}$ and is given by $\underline{\text{Hom}}(a, b) = a^\vee \otimes b$. In our case, as the internal Hom of $\text{Pr}_{\mathcal{V}}^L$ is given by $\mathcal{V}$ -linear colimit-preserving functors, we have

$$(13) \quad \text{Fun}^L(\text{Sh}(X), \text{Sh}(Y)) = \text{Sh}(X) \otimes \text{Sh}(Y) = \text{Sh}(X \times Y).$$

This identification can be given concretely as follows:

Definition 3.4.1. Let $F \in \text{Sh}(X)$ and $K \in \text{Sh}(X \times Y)$ and denote by p_1, p_2 the projections from $X \times Y$ to X and Y . Then, we use the notation

$$K \circ F := p_{2!}(K \otimes p_1^*F) \in \text{Sh}(Y).$$

This action is called convolution, and for these specific pairs (F, K) , we say that we convolve F with K . Similarly, let $L \in \text{Sh}(Y \times Z)$, and denote by q_{12}, q_{23}, q_{13} the projections from $X \times Y \times Z$ to $X \times Y$, $Y \times Z$, and $X \times Z$. Then, we use the notation

$$L \circ K := q_{13!}(q_{23}^*L \otimes q_{12}^*K) \in \text{Sh}(X \times Z).$$

Remark 3.4.2. As $!$ -pushforward can be thought of as a categorified fiberwise integration, the operation of convolving F with K can be thought of as a categorified version of convolving a function with a “kernel”, some nice function or measure on the product space. As a result, we sometimes refer to the sheaf K as a sheaf kernel, or simply a kernel, to emphasize the role it plays.

Proposition 3.4.3. *The identification (13) is given by*

$$\begin{aligned} \mathrm{Sh}(X \times Y) &\xrightarrow{\sim} \mathrm{Fun}^L(\mathrm{Sh}(X), \mathrm{Sh}(Y)) \\ K &\mapsto (F \mapsto K \circ F). \end{aligned}$$

Under this identification, composition of functors

$$\mathrm{Fun}^L(\mathrm{Sh}(Y), \mathrm{Sh}(Z)) \otimes \mathrm{Fun}^L(\mathrm{Sh}(X), \mathrm{Sh}(Y)) \rightarrow \mathrm{Fun}^L(\mathrm{Sh}(X), \mathrm{Sh}(Z)), (\beta, \alpha) \mapsto \beta \circ \alpha.$$

is identified with convolution of kernels

$$\mathrm{Sh}(Y \times Z) \otimes \mathrm{Sh}(X \times Y) \rightarrow \mathrm{Sh}(X \times Z), (L, K) \mapsto L \circ K.$$

Note that when $X = Y = Z$, convolution defines a monoidal structure on $\mathrm{Sh}(X \times X)$, which is identified with the natural monoidal structure on $\mathrm{End}^L(\mathrm{Sh}(X))$, the category of endofunctors of $\mathrm{Sh}(X)$.

Example 3.4.4. Let $f : X \rightarrow Y$ be a map in LCH, so we have $f_! \in \mathrm{Fun}^L(\mathrm{Sh}(X), \mathrm{Sh}(Y))$ and $f^* \in \mathrm{Fun}^L(\mathrm{Sh}(Y), \mathrm{Sh}(X))$. We claim that the corresponding kernels are given by 1_{Γ_f} and $1_{\Gamma_f^t}$ where $i_f : X \hookrightarrow \Gamma_f = \{(x, f(x)) | x \in X\}$ is the graph and Γ_f^t is the transport graph, the image of Γ_f under the coordinate swapping map $v : X \times Y \rightarrow Y \times X$ that is given by $v(x, y) = (y, x)$. As the computations are similar, we check the case for $f_!$. Recall first that, if $k : Z \hookrightarrow X$ is a locally closed embedding, then $F_Z := k_! k^*(F)$ is the same as $1_Z \otimes F$ where we use the simplified notation $1_Z := (1_X)_Z$. Indeed, this follows from the projection formula and the fact that $*$ -pullback is symmetric monoidal:

$$1_Z \otimes F = (k_! k^* 1_M) \otimes F = k_! (k^* 1_M \otimes k^* F) = k_! k^* (1_M \otimes F) = k_! k^* F.$$

But then, we note that, when we view Γ_f as a subset of $X \times Y$, we have $p_1 \circ i_f = \mathrm{id}_X$ and $p_2 \circ i_f = f$. Thus, for $F \in \mathrm{Sh}(X)$, we have

$$1_{\Gamma_f} \circ F = p_{2!}(1_{\Gamma_f} \otimes p_1^* F) = p_{2!} i_{f!} i_f^* p_1^* F = f_! F.$$

As a special case, when $X = Y$ and $f = \mathrm{id}_X$, we see that the identity functor $\mathrm{id}_{\mathrm{Sh}(X)}$ corresponds to the diagonal 1_Δ .

According to Proposition 3.4.3, having an action by isotopies can be reduced to constructing sheaf kernels that play well with the geometry. To state the main result, we first recall some basic notions from symplectic geometry.

Definition 3.4.5.

- (1) A symplectic manifold is a pair (X, ω) where X is a $2n$ -dimensional C^∞ manifold for some $n \in \mathbb{N}$ and $\omega \in \Omega^2(X)$ is a closed 2-form that is non-degenerate. This means that the natural pairing

$$\begin{aligned} TM &\rightarrow T^*M \\ v &\mapsto \omega(v, -) \end{aligned}$$

induces a vector bundle isomorphism $TM \cong T^*M$.

- (2) A (C^∞) isotopy $\Phi : X \times I \rightarrow X$ is a C^∞ map such that I is an interval containing 0 and, if we set $\varphi_t := \Phi(-, t)$ for $t \in I$, then $\varphi_t : X \rightarrow X$ is a diffeomorphism for all $t \in I$. An isotopy is said to be symplectic if, furthermore, $\varphi_t^* \omega = \omega$ for all $t \in I$; in other words, the family of diffeomorphisms is, in fact, a family of symplectomorphisms.

- (3) An exact symplectic manifold is a pair (X, λ) , where $\lambda \in \Omega^1(X)$ is such that $(X, d\lambda)$ is a symplectic manifold. The one-form is usually referred to as the Liouville form.
- (4) A symplectic isotopy $\Phi : X \times I \rightarrow X$ on an exact symplectic manifold (X, λ) is homogeneous if $\varphi_t^* \lambda = \lambda$.

Let (X, ω) be a symplectic manifold. The non-degeneracy condition implies that pairing with ω induces a vector space isomorphism

$$(14) \quad \begin{aligned} \mathfrak{X}(X) &\xrightarrow{\sim} \Omega^1(X) \\ v &\mapsto \iota_v \omega \end{aligned}$$

between vector fields and 1-forms. Let φ_t be a symplectic isotopy, and denote its associated (possibly time-dependent) vector field by v_t . When differentiating the equality $\varphi_t^* \omega = \omega$, Cartan's formula implies that

$$\frac{\partial}{\partial t} \varphi_t^* \omega = \varphi_t^* (\mathcal{L}_{v_t} \omega) = \varphi_t^* (d\iota_{v_t} \omega + \iota_{v_t} d\omega) = 0$$

or, $d(\iota_{v_t} \omega) = 0$, i.e., we see that the associated (possibly time-dependent) 1-form $\iota_{v_t} \omega$ is a closed 1-form.

Definition 3.4.6. In the situation as above, the symplectic isotopy φ_t is said to be a Hamiltonian isotopy if $\iota_{v_t} \omega$ is furthermore exact. A (possibly time-dependent) function H_t is said to be the Hamiltonian function or simply Hamiltonian if, for all $t \in I$, $dH_t = \iota_{v_t} \omega$.

Example 3.4.7. Let M be a C^∞ manifold. Its cotangent bundle T^*M admits a Liouville form $\lambda \in \Omega^1(T^*M) = \text{Hom}(TM, \mathbb{R}_M)$ that is given by

$$\lambda_{(x, \xi)}(\tilde{v}) = \xi(d\pi_{(x, \xi)}(\tilde{v})).$$

Here, we denote by $\pi : T^*M \rightarrow M$ the projection, and we use (x, ξ) to denote points in T^*M . In particular, a vector $\tilde{v} \in T_{(x, \xi)}M$ can be paired with ξ after being pushed to $T_x M$ by $d\pi_{(x, \xi)}$. This Liouville form is canonical in the sense that, if $f : M \rightarrow M$ is a diffeomorphism, then the induced diffeomorphism on T^*M preserves λ . One way to see that $d\lambda$ is indeed symplectic is by picking a chart of local coordinates $x = (x_1, \dots, x_n)$ on M , so it comes with cotangent coordinates $\xi = (\xi_1, \dots, \xi_n)$. Under these coordinates, we have $\lambda = \xi_1 dx_1 + \dots + \xi_n dx_n$, whose differential is clearly non-degenerate as it sends $\frac{\partial}{\partial x_k}$ to $-d\xi_k$ and $\frac{\partial}{\partial \xi_k}$ to dx_k .

Let (X, ω) be a symplectic manifold. The non-degeneracy condition implies that pairing with ω induces a vector space isomorphism

$$(15) \quad \begin{aligned} \mathfrak{X}(X) &\xrightarrow{\sim} \Omega^1(X) \\ v &\mapsto \iota_v \omega \end{aligned}$$

between vector fields and 1-forms. Let φ_t be a symplectic isotopy, and denote its associated (possibly time-dependent) vector field by v_t . When differentiating the equality $\varphi_t^* \omega = \omega$, Cartan's formula implies that

$$\frac{\partial}{\partial t} \varphi_t^* \omega = \varphi_t^* (\mathcal{L}_{v_t} \omega) = \varphi_t^* (d\iota_{v_t} \omega + \iota_{v_t} d\omega) = 0$$

or, $d(\iota_{v_t} \omega) = 0$, i.e., we see that the associated (possibly time-dependent) 1-form $\iota_{v_t} \omega$ is a closed 1-form, for each fixed t .

Definition 3.4.8. In the situation as above, the symplectic isotopy φ_t is said to be a Hamiltonian isotopy if $\iota_{v_t}\omega$ is furthermore exact. A (possibly time-dependent) function H_t is said to be the Hamiltonian function or simply Hamiltonian if, for all $t \in I$, $dH_t = \iota_{v_t}\omega$. Similarly, we call a time-dependent vector field v_t symplectic if $\iota_{v_t}\omega$ is closed for all t , and Hamiltonian if $\iota_{v_t}\omega$ is exact for all t .

Example 3.4.9. Let M be a C^∞ manifold. Its cotangent bundle T^*M admits a Liouville form $\lambda \in \Omega^1(T^*M) = \text{Hom}(TM, \mathbb{R}_M)$ that is given by

$$\lambda_{(x,\xi)}(\tilde{v}) = \xi(d\pi_{(x,\xi)}(\tilde{v})).$$

Here, we denote by $\pi : T^*M \rightarrow M$ the projection, and we use (x, ξ) to denote points in T^*M . In particular, a vector $\tilde{v} \in T_{(x,\xi)}M$ can be paired with ξ after being pushed to T_xM by $d\pi_{(x,\xi)}$. This Liouville form is canonical in the sense that, if $f : M \rightarrow M$ is a diffeomorphism, then the induced diffeomorphism on T^*M preserves λ . One way to see that $d\lambda$ is indeed symplectic is by picking a chart of local coordinates $x = (x_1, \dots, x_n)$ on M , so it comes with cotangent coordinates $\xi = (\xi_1, \dots, \xi_n)$. Under these coordinates, we have $\lambda = \xi_1 dx_1 + \dots + \xi_n dx_n$, whose differential is clearly non-degenerate as it sends $\frac{\partial}{\partial x_k}$ to $-d\xi_k$ and $\frac{\partial}{\partial \xi_k}$ to dx_k .

Example 3.4.10. The previous example shows that $\mathbb{R}^2 = T^*\mathbb{R}^1$ with $\lambda = \xi dx$ is an exact symplectic manifold. To obtain a simple example of a symplectic manifold that does not admit an exact structure, one notices that the action of $\mathbb{Z}^2$ on $\mathbb{R}^2$ by $(n, m) \cdot (x, \xi) = (x + n, \xi + m)$ preserves the symplectic form, so it descends to a symplectic form $\bar{\omega}$ on T^2 . In fact, $\bar{\omega}$ is the standard volume form on T^2 , which implies that it is not exact.

Example 3.4.11. Consider the symplectic manifold $\dot{T}^*S^1$ viewed as a cylinder $S^1 \times \mathbb{R}^1$ with its zero section $S^1 \times \{0\}$ removed. The isotopy $\Phi : \dot{T}^*S^1 \times \mathbb{R} \rightarrow \dot{T}^*S^1$ that is given by

$$\Phi(\theta, \xi, t) = \begin{cases} (\theta + t, \xi), & \xi > 0, \\ (\theta - t, \xi), & \xi < 0, \end{cases}$$

is homogeneous. Indeed, for fixed $t \in \mathbb{R}$, we have $\varphi_t^*(\xi d\theta) = \xi d(\theta + t) = \xi d\theta$. On the other hand, a similar computation implies that the isotopy $\Psi : \dot{T}^*S^1 \times [0, \infty) \rightarrow \dot{T}^*S^1$ that is given by

$$\Psi(\theta, \xi, t) = \begin{cases} (\theta, \xi + t), & \xi > 0, \\ (\theta, \xi - t), & \xi < 0, \end{cases}$$

is not homogeneous. However, one can compute that both isotopies are Hamiltonian with $-|\xi|$ being a Hamiltonian for Φ and $\text{sgn}(\xi)\theta$ being a Hamiltonian for Ψ . See also Remark 3.4.14 for more discussion.

Remark 3.4.12. In the case of $\dot{T}^*M$, a symplectic isotopy Φ is homogeneous if and only if it satisfies the simple condition that

$$\Phi(x, r\xi, t) = r\Phi(x, \xi, t), \quad \text{for all } r \in \mathbb{R}^+$$

where the group $\mathbb{R}^+$ of positive real numbers acts on $\dot{T}^*M$ via the vector bundle structure of T^*M . In general, an exact symplectic manifold (X, ω) admits a Liouville vector field X , corresponding to λ through (15), i.e., $\iota_X\omega = \lambda$. Thus, preserving λ is equivalent to preserving X , which, in turn, is equivalent to preserving the corresponding flow if X is integrable. In our case, the Liouville vector field is the Euler vector field, locally given by $X = \sum_{k=1}^n \xi_k \frac{\partial}{\partial \xi_k}$ and the flow, starting at (x, ξ) , is given by $(t \mapsto (x, e^t \xi))$.

Notation 3.4.13. Let $\Phi : \dot{T}^*M \times I \rightarrow \dot{T}^*M$ be a homogeneous symplectic isotopy. The “movie” Λ_Φ of Φ is the subset of $\dot{T}^*M \times \dot{T}^*M \times I$ that is given by the formula

$$\{(x, -\xi, \varphi_t(x, \xi), t, -H_\Phi(x, \xi, t)) \mid (x, \xi) \in \dot{T}^*M, t \in I\}.$$

Here, we use the notation $(\dot{-})$ to mean the partial derivative with respect to t , and one can think of the last term, $H_\Phi(x, \xi, t) := \lambda_{(x, \xi)}(\dot{\varphi}_t(x, \xi))$, as the (minus of the) “energy” of Φ . (See the next Remark 3.4.14 for more details.) For a time- t slice, we will use the notation

$$\Gamma_{\varphi_t}^a := \{(x, -\xi, \varphi_t(x, \xi)) \mid (x, \xi) \in \dot{T}^*M\}$$

for the anti-graph of φ_t .

Remark 3.4.14. A priori, “the” Hamiltonian H_t associated to a Hamiltonian isotopy φ_t is only well-defined up to a function $c(t)$ depending only on t . However, when (X, ω) comes from an exact symplectic manifold (X, λ) and φ_t is a homogeneous symplectic isotopy, differentiating the equation $\varphi_t^*\lambda = \lambda$ will imply that $\iota_{v_t}\omega = \iota_{v_t}d\lambda = -d(\lambda(v_t))$. In fact, assuming the Liouville vector field integrates to an $\mathbb{R}$ -action, then the exact form induces a bijection

$$\begin{aligned} \{\text{symplectic vector fields}\} &\rightarrow C^\infty(X)^{\mathbb{R}} \\ v_t &\mapsto \lambda(v_t). \end{aligned}$$

Thus, the homogeneous Hamiltonian $H_\Phi = -\lambda(v_t)$ can be thought of as the “energy” of this isotopy.

We now state the main theorem for this section:

Theorem 3.4.15 ([28]). *Let $\Phi : \dot{T}^*M \times I \rightarrow \dot{T}^*M$ be a homogeneous symplectic isotopy. Then there exists a unique “ I -family” sheaf kernel $K(\Phi) \in \text{Sh}(M \times M \times I)$ such that*

- (1) $K(\Phi)|_0 = 1_\Delta$,
- (2) $\text{SS}(K(\Phi)) \subseteq \Lambda_\Phi \cup 0_{M \times M \times I}$.

The kernel $K(\Phi)$ is usually referred to as the Guillermou-Kashiwara-Schapira sheaf quantization³⁰ kernel, or simply as the GKS kernel, and the procedure is referred to as GKS sheaf quantization.

Notation 3.4.16. For $t \in I$, we will use the notation $i_t : M \times M \hookrightarrow M \times M \times I$ for the inclusion of the slice at t . Furthermore, for an I -family sheaf $K \in \text{Sh}(M \times M \times I)$, we will use the notation $K|_t := i_t^*K = K|_{M \times M \times \{t\}}$ to mean its restriction to this slice.

Lemma 3.4.17. *For $t \in I$, we have $\text{SS}(K(\Phi)|_t) \subseteq \Gamma_{\varphi_t}^a$.*

Proof. Fix $t \in I$. We note that, regardless of what the last term in Λ_Φ is, the points in Λ_Φ cannot have covectors non-zero only on the I -direction. Thus, condition (2) of Theorem 3.4.15 implies that $K(\Phi)$ is non-characteristic for i_t and (1) of Proposition 3.3.11 allows us to estimate the microsupport of $K(\Phi)|_t$ by Example 3.3.5: First restrict to the t -slice from the base, which in this case gives us

$$\{(x, -\xi, \varphi_t(x, \xi), -\lambda_{(x, \xi)}(\dot{\varphi}_t(x, \xi)) \mid (x, \xi) \in \dot{T}^*M\} \subseteq \dot{T}^*M \times \dot{T}^*M \times T_t^*I,$$

and then project out the cotangent direction on I , which leads to

$$\Gamma_{\varphi_t}^a = \{(x, -\xi, \varphi_t(x, \xi)) \mid (x, \xi) \in \dot{T}^*M\}.$$

□

³⁰Following the physical convention that a symplectic isotopy should “quantize” to a family of operators.




FIGURE 7. The picture of the slices of $K(\tilde{\Phi})$ in the case when $\tilde{\Phi}(x, \xi, t) = (x, \xi \pm t)$ on $\dot{T}^*\mathbb{R}$. The middle graph is the diagonal, the left is the picture for $t < 0$, and the right is the picture for $t > 0$. One notices that the two boundaries are both given by $y = x \pm t$, but the requirement from the cotangent direction determines whether the boundary should be closed or open.

To get a feeling of this theorem, let us consider an example.

Example 3.4.18. Consider the case when $M = \mathbb{R}^1$ and $\tilde{\Phi} : \dot{T}^*\mathbb{R}^1 \times \mathbb{R}^1 \rightarrow \dot{T}^*\mathbb{R}^1$ given by

$$\tilde{\Phi}(x, \xi, t) = \begin{cases} (x + t, \xi), & \xi > 0, \\ (x - t, \xi), & \xi < 0. \end{cases}$$

This is the lift of Φ from Example 3.4.11 to the universal cover. As $\dot{T}^*\mathbb{R}$ consists of two components, its anti-graph $\Gamma_{\tilde{\varphi}_t}^a$ is the union

$$\{(x, -\xi, x + t, \xi) | x \in \mathbb{R}^1, \xi > 0, t \in \mathbb{R}^1\} \cup \{(x, -\xi, x - t, \xi) | x \in \mathbb{R}^1, \xi < 0, t \in \mathbb{R}^1\}.$$

By the previous Lemma 3.4.17, $K(\tilde{\Phi})|_t$, the slice at t is microsupported in $\Gamma_{\tilde{\varphi}_t}^a$. When $t = 0$, the slice is the diagonal 1_Δ , so a reasonable guess is that, on the base, we should begin from (x, x) and, for example when $t > 0$, move up toward $(x, x + t)$ and down toward $(x, x - t)$, and then assign the correct boundary according to the direction toward which $(-\xi, \xi)$ is pointing. To be more concrete, consider $t = 1$ and $\xi = 1$. The point $(x, x, -1, 1)$ will be sent to $(x, x + 1, -1, 1)$. Similarly, the point $(x, x, 1, -1)$ will be sent to $(x, x - 1, 1, -1)$. Pictorially, we thus conclude that the time-1 slice $K(\tilde{\Phi})|_1$ should have support given by the region $\{|y - x| \leq t\}$ and the microsupport at infinity occurs only on the boundary as the outward conormal, indicating an open boundary, per Example 3.3.3. A similar discussion implies that the time-(-1) slice has the same support, with inward conormal on the boundary, so it has a closed boundary. A reasonable guess for the answer, up to shifting, which we visualize in Figure 7, is then

$$K(\tilde{\Phi}) = \begin{cases} 1_{\{x-t < y < x+t\}}[2], & t > 0, \\ 1_{\{x+t \leq y \leq x-t\}}, & t \leq 0. \end{cases}$$

To get the entire $K(\tilde{\Phi})$, essentially what one does is to notice that, if we allow t to vary, we obtain two sheaves $1_{\{x-t < y < x+t, t > 0\}}[1]$ and $1_{\{x+t \leq y \leq x-t, t \leq 0\}}$ on $\mathbb{R}_x^1 \times \mathbb{R}_y^1 \times \mathbb{R}_t^1$. One can compute that, away from $t = 0$, they give the correct microsupport and the claim is that $K(\tilde{\Phi})$ can be “glued” from them. More precisely, there exists a canonical map $\alpha : 1_{\{x+t \leq y \leq x-t, t \leq 0\}} \rightarrow 1_{\{x-t < y < x+t, t > 0\}}[2]$ and $K(\tilde{\Phi}) = \text{fib}(\alpha)$.

The main goal of this section is to explain the proof of Theorem 3.4.15. As the proof requires many steps, we first summarize them and indicate where each step is achieved. Before we begin, since we cannot assume the existence or uniqueness of $K(\Phi)$, we will use the following terminology:

Definition 3.4.19. Let Φ be a homogeneous isotopy on $\dot{T}^*M$ and $K \in \text{Sh}(M \times M \times I)$. We say that K is a GKS kernel for Φ if K satisfies conditions (1) and (2) of Theorem 3.4.15.

The first step is to establish the compatibility between convolution of GKS kernels and the composition of their underlying isotopies. Since isotopies are always invertible, this would imply that if K and L are GKS kernels of an isotopy Φ and its inverse Φ^{-1} , then both K and L are inverse to each other with respect to convolution. See Corollary 3.4.27. Notice that, even if Φ has a GKS kernel K , we have not said whether Φ^{-1} admits a GKS kernel. In fact, what we have shown is more subtle: if both K_1 and K_2 are GKS kernels for Φ , then any GKS kernel L for Φ^{-1} will be their inverse. But the same holds vice versa, and uniqueness of the convolution inverse then implies that K_1 and K_2 are in fact the same. Thus, to establish uniqueness, it suffices to show that if K is a GKS kernel for Φ , one can construct a GKS kernel for Φ^{-1} , which we temporarily denote by $K^{-1, \text{pre}}$, using the six-functor formalism. See Lemma 3.4.37 and Proposition 3.4.38. For the discussion of existence, we restrict ourselves to a much simpler case when M is a compact manifold and Φ admits a time-independent Hamiltonian. The next step is to use uniqueness and the compatibility to reduce an existence statement for a general Φ to a special case when Φ is somewhat “small”. See Lemma 3.4.39. Lastly, we will explain how this small case is done in Proposition 3.4.40.

Definition 3.4.20. Let $\Phi, \Psi : \dot{T}^*M \times I \rightarrow \dot{T}^*M$ be homogeneous symplectic isotopies. The composition isotopy $(\Psi \circ |_I \Phi) : \dot{T}^*M \times I \rightarrow \dot{T}^*M$ is given by $(\Psi \circ |_I \Phi)(x, \xi, t) = \psi_t(\varphi_t(x, \xi))$. Note that under this composition, there is a unit isotopy $\text{Id} : \dot{T}^*M \times I \rightarrow \dot{T}^*M$ given by $\text{Id}(x, \xi, t) = (x, \xi)$ and any Φ has an inverse isotopy Φ^{-1} given by $\Phi^{-1}(x, \xi, t) = \varphi_t^{-1}(x, \xi)$.

Lemma 3.4.21. *We have the following identities:*

- (1) $H_{\text{Id}} = 0$, in particular $\Lambda_{\text{Id}} = N^*(\Delta) \times 0_I$.
- (2) If Φ, Ψ are homogeneous symplectic isotopies, then

$$H_{\Psi \circ |_I \Phi}(x, \xi, t) = H_{\Psi}(\varphi_t(x, \xi), t) + H_{\Phi}(x, \xi, t).$$

- (3) $H_{\Phi^{-1}}(x, \xi, t) = -H_{\Phi}(\varphi_t^{-1}(x, \xi), t)$.

Proof. By the chain rule, we have

$$\frac{\partial}{\partial t}(\psi_t \circ \phi_t(x, \xi)) = \dot{\psi}_t(\varphi_t(x, \xi)) + d(\psi_t)_{\varphi_t(x, \xi)}(\dot{\varphi}_t(x, \xi)).$$

Thus, substituting these into $\lambda_{\psi_t \circ \varphi_t}(x, \xi)$, we obtain

$$H_{\Psi \circ |_I \Phi}(x, \xi, t) = \lambda_{\psi_t \circ \varphi_t(x, \xi)}(\dot{\psi}_t(\varphi_t(x, \xi)) + (\psi_t^* \lambda)_{\varphi_t(x, \xi)}(\dot{\varphi}_t(x, \xi))).$$

The second identity then follows from the definition $H_{\Phi}(x, \xi, t) = \lambda_{\varphi_t(x, \xi)}(\dot{\varphi}_t(x, \xi))$ and the fact that $\varphi_t^* \lambda = \lambda$. To obtain the third identity, one uses the equality $\Phi^{-1} \circ \Phi = \text{Id}$. $\square$

Definition 3.4.22. For $K, L \in \text{Sh}(M \times M \times I)$, we denote by $q_{12I}, q_{23I}, q_{13I}$ the projections from $X \times Y \times Z \times I$ to $X \times Y \times I$, $Y \times Z \times I$, and $X \times Z \times I$. Then, we use the notation

$$L \circ |_I K := q_{13I!}(q_{23I}^* L \otimes q_{12I}^* K) \in \text{Sh}(X \times Z \times I),$$

and call it the convolution relative to I . The unit of the convolution is $1_{\Delta \times I}$.

We note that base change implies the following compatibility result:

Lemma 3.4.23. *Let $K, L \in \text{Sh}(M \times M \times I)$ be two I -family sheaf kernels. Then, for any $t \in I$, we have the equality*

$$(L \circ |_I K)|_t = L|_t \circ K|_t.$$

We begin our study of the GKS kernel by noticing that the slices of $K(\tilde{\Phi})$ from Example 3.4.18 have the feature that the two projections $p_1, p_2 : \mathbb{R}^2 \rightarrow \mathbb{R}^1$ are both proper. In fact, this holds true in general.

Lemma 3.4.24. *Fix a symplectic isotopy Φ . Let $K \in \text{Sh}(M \times M \times I)$ be a kernel satisfying conditions (1) and (2) in Theorem 3.4.15 for Φ . Then the two projections $p_{1I}, p_{2I} : M \times M \times I \rightarrow M \times I$ are proper on the support $\text{supp}(K)$.*

Proof. We note that the intersection $\overline{\Lambda_\Phi} \cap 0_{M \times M \times I}$ is proper with respect to p_{1I} and p_{2I} . By (2) in Theorem 3.4.15 and Example 3.2.10, the sheaf K is a local system away from $\overline{\Lambda_\Phi} \cap 0_{M \times M \times I}$. But (1) in Theorem 3.4.15 then implies that its stalks there must be 0. $\square$

This observation allows us to use Proposition 3.3.11 to estimate the microsupport for pushforward, which is part of the building block of convolution. We thus have:

Lemma 3.4.25. *Let $\Phi, \Psi : \dot{T}^*M \times I \rightarrow \dot{T}^*M$ be two homogeneous isotopies indexed by the same interval I . If K and L are GKS kernels for Φ and Ψ , then $L \circ |_I K$ is a GKS kernel for the composition isotopy $\Psi \circ |_I \Phi$.*

Proof. Condition (1) follows from Lemma 3.4.23. For condition (2), we will actually only prove the simpler case for the slice at t , i.e., $L|_t \circ K|_t$ is microsupported in $\Gamma_{\psi_t \circ \varphi_t}^a$; this case requires the same techniques from microlocal sheaf theory, and the extra complexity comes from further tracing the energy term via Lemma 3.4.21 for each step. To see this, we see that since $\text{SS}(q_{23}^* L|_t) = 0_M \times \text{SS}(L|_t)$ and $\text{SS}(q_{12}^* K|_t) = \text{SS}(K|_t) \times 0_M$, the condition for estimating the microsupport of the product, $\text{SS}(q_{23}^* L|_t)^a \cap \text{SS}(q_{12}^* K|_t) \subseteq 0_{M \times M \times M}$, is satisfied and Proposition 3.3.18 implies that the tensor product $q_{23}^* L|_t \otimes q_{12}^* K|_t$ is microsupported in the fiberwise sum $S := \text{SS}(q_{23}^* L|_t) + \text{SS}(q_{12}^* K|_t)$. Furthermore, as observed in the previous Lemma 3.4.24, pushing along q_{13} is proper and we can apply Proposition 3.3.11. A point (x, ξ, z, ζ) in $\text{SS}(L|_t \circ K|_t)$ is thus in the set $(q_{13})_\pi([dq_{13}^\vee]^{-1}(S))$. Example 3.3.5 then implies that there is some y such that $(x, \xi, y, 0, z, \zeta)$ in the fiberwise sum S . Because the only non-trivial component that both sets share is the (y, η) -component, this is equivalent to having some (y, η) such that $\varphi_t(x, \xi) + (y, -\eta) = 0$ and $(z, \zeta) = \psi_t(y, \eta)$. Combining the two conditions, we see that $(z, \zeta) = \psi_t(y, \eta) = \psi_t(\varphi_t(x, \xi))$, as desired; one also sees that this is why we need the “ $-$ ” before the first ξ in $\Gamma_{\varphi_t}^a$. $\square$

Remark 3.4.26. The above Lemma 3.4.25 generalizes, with similar noncharacteristic and proper assumptions, to the compatibility between convolution and composition of relations. See for example [28, (1.12)].

Corollary 3.4.27. *If K and L are GKS kernels for Φ and Φ^{-1} , then K and L are inverse to each other. That is, $K \circ |_I L = L \circ |_I K = 1_{\Delta \times I}$.*

Proof. Lemma 3.4.25 implies that the compositions $K \circ |_I L$ and $L \circ |_I K$ are both GKS kernels of the identity isotopy, so the corollary boils down to proving uniqueness for this case. By Lemma 3.4.21, its movie is $N^*(\Delta) \times 0_M$ so the next Lemma 3.4.28 implies that they are both of the form $H|_0 \boxtimes 1_I$, and the requirement at $t = 0$ pins down $H|_0$ as 1_Δ . $\square$

Lemma 3.4.28. *Let N be a C^∞ manifold and $p : N \times I \rightarrow N$ be the projection to the first component. Then p^* induces an equivalence $p^* : \mathrm{Sh}(N) \xrightarrow{\sim} \mathrm{Sh}_{T^*N \times 0_I}(N \times I)$. In particular, if $\mathrm{SS}(H) \subseteq T^*N \times 0_I$, then $H = (H|_{N \times \{0\}}) \boxtimes 1_I$.*

Sketch of the proof. The first statement would imply that, for $H \in \mathrm{Sh}_{T^*N \times 0_I}$, the counit $(p_*H) \boxtimes 1_I = p^*p_*H \xrightarrow{\sim} H$ is an isomorphism. But restricting to $N \times \{0\}$ then implies that $H|_{N \times \{0\}} = p_*H$. To see the equivalence, we need to check that the unit $p^*p_*H \rightarrow H$ for $H \in \mathrm{Sh}_{T^*N \times 0_I}(N \times I)$ and the counit $F \rightarrow p_*p^*F$ are equivalences.

We focus on the unit, as the steps for the counit are similar but simpler. The claim is that one can replace I by the closed interval $[-1, 1]$, and then reduce to the case when $N = \{*\}$ is a point, in which case, we will be proving that $\mathrm{Loc}([-1, 1]) \cong \mathcal{V}$, which can be done by the homotopical method explained in [36, Section 2.7]. Before the reduction, we notice that H is noncharacteristic for the inclusion $i_x : \{x\} \times I \hookrightarrow N \times I$. Proposition 3.3.11 and Example 3.2.10 then show that H is fiberwise a local system, i.e., $H|_{\{x\} \times I} \in \mathrm{Loc}(I)$, so it suffices to show that the unit is an isomorphism for such sheaves.

To perform the reduction, take an increasing family of closed intervals $\overline{I}_n$ such that $\bigcup_{n \in \mathbb{Z}} \overline{I}_n = I$. Denote by $f_n : N \times \overline{I}_n \hookrightarrow N \times I$ the closed inclusion. One can check that the unit $H \rightarrow f_{n,*}f_n^*H$ induces an isomorphism $H = \lim_{n \rightarrow \infty} f_{n,*}f_n^*H$. Set $\overline{p}_n := p \circ f_n$. Since both $p^* = p^![-1]$ and p_* are right adjoints, we have $p^*p_*H = \lim_{n \rightarrow \infty} \overline{p}_n^*\overline{p}_{n*}H|_{N \times \overline{I}_n}$ and we reduce to the case of $\overline{p} : N \times [-1, 1] \rightarrow N$. For this case, as $\overline{p}$ is proper, one can check at stalks by base change, so we reduce to the case when $N = \{*\}$. $\square$

With Corollary 3.4.27, our goal is now to show that if K is a GKS kernel for Φ , then there is a way to produce a candidate kernel $K^{-1, \mathrm{pre}}$ that is potentially a GKS kernel for Φ^{-1} . To get a formula like this, we make the following consideration: For $K \in \mathrm{Sh}(X \times Y)$, the functor $K \circ (-) : \mathrm{Sh}(X) \rightarrow \mathrm{Sh}(Y)$ has a right adjoint

$$\mathcal{H}\mathrm{om}^\circ(K, G) := p_{1*} \mathcal{H}\mathrm{om}(K, p_2^!G).$$

Now, if $K \circ (-)$ happens to be an equivalence, it in particular implies that this right adjoint is also an equivalence and in particular is colimit-preserving so, by Proposition 3.4.3. So the general question is how we can obtain an expression $\mathcal{H}\mathrm{om}^\circ(K, -) = K^R \circ (-)$, for some kernel $K^R \in \mathrm{Sh}(Y \times X)$ when we expect the right adjoint to be colimit-preserving. In other words, the main task is to turn a Hom into a tensor.

At the level of algebra, we have seen such a situation: We recall that, in a symmetric monoidal category $\mathcal{V}$, an internal hom object $\underline{\mathrm{Hom}}(x, y) \in \mathcal{V}$, for $x, y \in \mathcal{V}$, is classified by the equation

$$\mathrm{Hom}_{\mathcal{V}}(z, \underline{\mathrm{Hom}}(x, y)) = \mathrm{Hom}_{\mathcal{V}}(z \otimes x, y)$$

for $z \in \mathcal{V}$. If $\mathcal{V} \in \mathrm{CAlg}(\mathrm{Pr}_{st}^{\mathrm{L}})$ is a symmetric monoidal presentable category, then $\underline{\mathrm{Hom}}(x, y)$ always exists by the adjoint functor theorem, Theorem 2.2.6, since $\underline{\mathrm{Hom}}(x, -)$ is the right adjoint of $x \otimes (-)$, which is assumed to be colimit-preserving. For any $x, y \in \mathcal{V}$, we then have the canonical map, considered in (11)

$$\underline{\mathrm{Hom}}(x, 1) \otimes y \rightarrow \underline{\mathrm{Hom}}(x, y),$$

which we know is an equivalence, for example, when x is dualizable. Note that a necessary condition for this map to be an equivalence is that the expression be colimit-preserving in y on the left-hand side and limit-preserving in y on the right-hand side.

Example 3.4.29. A convenient setup is the following: A symmetric monoidal $\mathcal{V}_0 \in \text{CAlg}(\text{st})$ is said to be rigid if all $x \in \mathcal{V}_0$ are dualizable. In this case, the symmetric monoidal structure $-\otimes- : \mathcal{V}_0 \otimes \mathcal{V}_0 \rightarrow \mathcal{V}_0$ induces a symmetric monoidal structure $\mathcal{V} \in \text{CAlg}(\text{Pr}_{\omega, \text{st}}^{\text{L}})$ on its Ind-completion. In this case, all compact objects in $\mathcal{V}$ are dualizable. The category of compact objects in spectra, Sp^ω , and the category $\text{Perf}(\mathbb{Z})$ of perfect complexes of abelian groups, are standard examples of rigid categories.

Now, we want to consider the generalization to sheaves. In addition to the algebraic condition, we also need some control on the geometry.

Definition 3.4.30. Let $F \in \text{Sh}(X)$ be a sheaf.

- (1) We say F has perfect stalks if F_x are dualizable in $\mathcal{V}$ for all $x \in X$.
- (2) We say F is constructible if there exists a Whitney stratification $X = \coprod_{\alpha \in s} X_\alpha$, a decomposition of X into locally finite, locally closed subsets X_α satisfying the Whitney conditions, e.g. [36, Exercise VIII. I 2], such that the restriction of F to each stratum, $F|_{X_\alpha}$, lies in $\text{Loc}(X_\alpha)$.

Proposition 3.4.31 ([36, Proposition 3.4.3, Proposition 5.4.14]). *Let $G, F \in \text{Sh}(X)$ be sheaves and assume G is constructible with perfect stalks. Then,*

- (1) *The canonical map $G \rightarrow D_M D_M(G)$ is an equivalence, where $D_M(G) := \mathcal{H}\text{om}(G, \omega_X)$ is contravariant Verdier duality.*
- (2) *If $\text{SS}(F) \cap \text{SS}(G) \subseteq 0_M$, then the canonical map*

$$\mathcal{H}\text{om}(G, 1_M) \otimes F \rightarrow \mathcal{H}\text{om}(G, F)$$

is an equivalence.

Remark 3.4.32. The above Proposition 3.4.31 holds more generally for sheaves that are “cohomologically” constructible with perfect stalks and costalks. (Note that we are following the modern convention, e.g., in [50], of using the term “constructibility” to refer only to the pointwise geometry, separating the perfectness property from the term. Classically, e.g., in [36], a sheaf with constructible geometry but without perfectness on stalks and costalks is referred to as “quasi-constructible”.)

Remark 3.4.33. The (non-invertible) contravariant Verdier duality is obtained by the covariant Verdier duality in Theorem 2.6.6 by composing with the naive linear dual. More precisely, for a sheaf F , its covariant Verdier dual is the cosheaf given by the assignment $(U \mapsto \Gamma_a(U; F))$ for $U \subseteq X$ open. Taking $\text{Hom}(-, 1_{\mathcal{V}})$, we see that

$$\text{Hom}(\Gamma_a(U; F), 1_{\mathcal{V}}) = \text{Hom}(F|_U, \omega_U) = \Gamma(U; \mathcal{H}\text{om}(F, \omega_X)).$$

Proposition 3.4.34. *Let X and Y be manifolds and let $K \in \text{Sh}(X \times Y)$ be a sheaf kernel. Assume that*

- (1) *The projection to X , $p_X : X \times Y \rightarrow X$, is proper on $\text{supp}(K)$,*
- (2) *K is constructible with perfect stalks, and*
- (3) *$\text{SS}(K) \cap 0_X \times T^*Y \subseteq 0_{X \times Y}$.*

Then, $K \circ (-)$ has a colimit-preserving right adjoint given by $K^R \circ (-)$ where

$$K^R := \mathcal{H}\text{om}(v^* K, 1_Y \boxtimes \omega_X)$$

where $v : X \times Y \cong Y \times X$ is the coordinate swapping map $v(x, y) = (y, x)$. ³¹

³¹Applying v is not essential but it is psychologically easier to keep track of the direction of the functor when $X = Y$.

Proof. Let $G \in \text{Sh}(Y)$. Our goal is to turn $p_{X*} \mathcal{H}\text{om}(K, p_Y^! G)$ into a convolution. By Proposition 3.3.11, $\text{SS}(p_Y^! G)$ is microsupported in $0_X \times \text{SS}(G)$ and $p_Y^! G = p_Y^* G \otimes p_Y^! 1_Y$. Assumptions (2) and (3) then imply that Proposition 3.4.31 is applicable and $\mathcal{H}\text{om}(K, p_Y^! G) = \mathcal{H}\text{om}(K, 1_{X \times Y}) \otimes p_Y^! G$. To get the correct form for convolution, recall that we saw in the proof of Lemma 3.3.7 that $p_Y^! 1_Y = p_X^* \omega_X = \omega_X \boxtimes 1_Y$. Lastly, Assumption (1) implies that

$$p_{X*} \mathcal{H}\text{om}(K, p_Y^! G) = p_{X!} (\mathcal{H}\text{om}(K, \omega_X \boxtimes 1_Y) \otimes p_Y^* G) = p_{X!} (\mathcal{H}\text{om}(v^* K, 1_Y \boxtimes \omega_X) \otimes p_Y^* G),$$

as desired. $\square$

We remark that we do not actually know if a GKS kernel K for Φ satisfies the conditions in Proposition 3.4.34. Nevertheless, the left expression is valid for any sheaf kernel and we consider the object regardless.

Notation 3.4.35. Let $K \in \text{Sh}(M \times M)$ be a sheaf kernel. Denote by $v : M \times M \rightarrow M \times M$ the coordinate swapping map $v(x, y) = (y, x)$. The candidate convolution inverse of K is given by

$$K^{-1, \text{pre}} := \mathcal{H}\text{om}(v^* K, 1_M \boxtimes \omega_M).$$

For the I -family case, let $K \in \text{Sh}(M \times M \times I)$ be an I -family kernel. We abuse notation and also use $K^{-1, \text{pre}}$ to denote the candidate convolution inverse with respect to $\circ|_I$

$$K^{-1, \text{pre}} := \mathcal{H}\text{om}((v \times \text{id}_I)^* K, 1_M \boxtimes \omega_M \boxtimes 1_I).$$

Example 3.4.36. As a sanity check, we compute the case when $K = 1_\Delta$:

$$1_\Delta^{-1, \text{pre}} := \mathcal{H}\text{om}(1_\Delta, 1_M \boxtimes \omega_M) = \Delta_* \Delta^! p_2^! 1_M = \Delta_* 1_M = 1_\Delta.$$

That is, $1_\Delta^{-1, \text{pre}} = 1_\Delta^{-1} = 1_\Delta$ as it corresponds to the identity functor. Note that, despite the notation suggests $K^{-1, \text{pre}}$, the kernel can be non-invertible. For example, there is the trivial example of $0^{-1, \text{pre}} = 0$. A more interesting example is when $K = \bigoplus_{n \in \mathbb{Z}} 1_\Delta$. In this case,

$$K^{-1, \text{pre}} = \mathcal{H}\text{om}(\bigoplus_{n \in \mathbb{Z}} 1_\Delta, 1_M \boxtimes \omega_M) = \prod_{n \in \mathbb{Z}} \mathcal{H}\text{om}(\bigoplus_{n \in \mathbb{Z}} 1_\Delta, 1_M \boxtimes \omega_M) = \prod_{n \in \mathbb{Z}} 1_\Delta.$$

Lemma 3.4.37. Let $\Phi : \dot{T}^* M \times I \rightarrow \dot{T}^* M$ be a homogeneous symplectic isotopy. If K is a GKS kernel for Φ , then $K^{-1, \text{pre}}$ is a GKS kernel for Φ^{-1} .

Proof. By Proposition 3.3.18, $\dot{\text{SS}}(K^{-1, \text{pre}}) \subseteq \dot{\text{SS}}((v \times \text{id})^* K)^a$ so it consists of points of the form $(-\varphi_t(x, \xi), x, \xi, t, H_\Phi(x, \xi, t))$. Applying the change of coordinates $(y, \eta) = \varphi_t(x, \xi)$, we see that $(x, \xi) = \varphi_t^{-1}(y, \eta)$ so those points are of the form $(y, -\eta, \varphi_t^{-1}(x, \xi), t, -H_{\Phi^{-1}}(y, \eta, t))$, which implies that $\dot{\text{SS}}(K^{-1, \text{pre}}) \subseteq \Lambda_{\Phi^{-1}}$ as desired. Here, we use the equation $H_{\Phi^{-1}}(x, \xi, t) = -H_\Phi(\varphi_t^{-1}(x, \xi), t)$ from Lemma 3.4.21. To see $K^{-1, \text{pre}}|_0 = 1_\Delta$, we need to commute i_0^* with $\mathcal{H}\text{om}$. We notice that for any $H, L \in \text{Sh}(M \times M \times I)$, we have the following commuting diagram

$$\begin{array}{ccc} i_0^*(\mathcal{H}\text{om}(L, H)) \otimes \omega_{i_0} & \longrightarrow & \mathcal{H}\text{om}(i_0^* L, i_0^*(H) \otimes \omega_{i_0}) \\ \downarrow & & \downarrow \\ i_0^!(\mathcal{H}\text{om}(L, H)) & \xrightarrow{\sim} & \mathcal{H}\text{om}(i_0^* L, i_0^! H) \end{array}$$

so it suffices to show that $i_0^*(-) \otimes \omega_{i_0} \rightarrow i_0^!(-)$ is an isomorphism in our case and that ω_{i_0} is invertible. The former follows from Proposition 3.3.11 and the computation $\dot{\text{SS}}(K^{-1, \text{pre}}) \subseteq \Lambda_{\Phi^{-1}}$ above, and the latter follows from Lemma 3.3.13. $\square$

Proposition 3.4.38. *Let $\Phi : \dot{T}^*M \times I \rightarrow \dot{T}^*M$ be a homogeneous symplectic isotopy. GKS kernels for Φ , if they exist, are unique, so the notation $K(\Phi)$ is well-defined. Furthermore, if $K(\Phi)$ exists, then $K(\Phi^{-1})$ also exists and is given by $K(\Phi)^{-1} = K(\Phi)^{-1, \text{pre}}$.*

Proof. Let K_1, K_2 be GKS kernels for Φ . By the above Lemma 3.4.37, $K_1^{-1, \text{pre}}$ is a GKS kernel for Φ^{-1} . Corollary 3.4.27 then implies that $K_1^{-1, \text{pre}}$ is both the inverse of K_1 and K_2 . Thus,

$$K_2 = K_2 \circ |_I 1_{\Delta \times \{0\}} = (K_2 \circ |_I K_1^{-1, \text{pre}}) \circ |_I K_1 = 1_{\Delta \times \{0\}} \circ |_I K_1 = K_1.$$

The rest also follows from Corollary 3.4.27 and Lemma 3.4.37. $\square$

We now explain how we can use the uniqueness property to reduce the proof of existence. The general strategy is to use the fact that $K(\Psi \circ |_I \Phi) = K(\Psi) \circ |_I K(\Phi)$ provided that both $K(\Psi)$ and $K(\Phi)$ exist, and try to decompose a complicated Φ into simpler pieces.

We will restrict ourselves to the simplified case when M is compact and H_Φ is time-independent. Since $H_\Phi(x, \xi, t)$ is homogeneous, i.e., $H_\Phi(x, r\xi, t) = rH_\Phi(x, \xi, t)$ for $r \in \mathbb{R}^+$, it descends to a function $\overline{H}_\Phi([x, \xi], t)$ on the cosphere bundle $S^*M := \dot{T}^*M/\mathbb{R}^+$. Fix a Riemannian metric g , then we can identify $S^*M \subseteq \dot{T}^*M$ as the subset consisting of unit covectors, so $H_\Phi(x, \xi, t) = |\xi| \overline{H}_\Phi([x, \xi], t)$. Since M is compact and $H_\Phi(x, \xi, t) = H_\Phi(x, \xi)$ is time-independent, there exists some $C > 0$ such that $C + \min_{[x, \xi] \in S^*M} \overline{H}_\Phi([x, \xi]) > 0$. Non-degeneracy of $d\lambda$ implies that there is a Hamiltonian vector field v on $\dot{T}^*M$ corresponding to the function $C|\xi|$ and it integrates to a homogeneous symplectic isotopy $\Psi : \dot{T}^*M \times I \rightarrow \dot{T}^*M$, since M is compact (and v is homogeneous). Our choice of C implies that $\Psi \circ \Phi$ has positive energy. Since $\Phi = \Psi^{-1} \circ |_I \Psi \circ |_I \Phi$, Lemma 3.4.25, Lemma 3.4.21, and Proposition 3.4.38 imply

Lemma 3.4.39. *When M is compact, if the GKS kernel $K(\Phi)$ exists for any Φ such that $H_\Phi > 0$, then $K(\Psi)$ exists for any time-dependent symplectic homogeneous isotopy Ψ . In fact, given any $\epsilon > 0$, it suffices to prove existence for those Φ such that $\epsilon|\xi| > H_\Phi > 0$.*

Proof. To see the “in fact” part, note that for time-independent Φ , if we set $\Phi_n(x, \xi, t) = \Phi(x, \xi, t/n)$, then $H_{\Phi_n} = \frac{1}{n}H_\Phi$. $\square$

Proposition 3.4.40. *When M is compact, $K(\Phi)$ exists for time-dependent homogeneous symplectic isotopy Φ .*

Sketch of proof. The claim, the details of which we will skip, is that the construction in Example 3.4.18 essentially works. More precisely, denote by d_g the distance function on M induced by g . For any $\epsilon \ll 1$, much smaller than the injectivity radius, so that the set $\{(x, y) \mid d_g(x, y) < 2\epsilon\}$ is a tubular neighborhood of 1_Δ , the time- t slice for small negative time $t < 0$ is given by $1_{\{d_g(x, y) \leq t\epsilon\}}$. The picture for small positive time $t > 0$ is likewise $\mathcal{H}\text{om}(1_{\{d_g(x, y) \leq t\epsilon\}}, 1_M \boxtimes \omega_M)$ and one can organize them into the GKS kernel $K(\Psi_\epsilon)$ associated to the Hamiltonian $\epsilon|\xi|$. For a general small enough strictly positive Φ such that $\epsilon|\xi| > H_\Phi > \epsilon'|\xi| > 0$, $K(\Phi)$ will be given by an interpolation between $K(\Phi_\epsilon)$ and $K(\Phi_{\epsilon'})$. $\square$

Example 3.4.41. Consider the case when $M = S^1$ and the isotopy is given by

$$\Phi(\theta, \xi, t) = \begin{cases} (\theta + t, \xi), & \xi > 0, \\ (\theta - t, \xi), & \xi < 0, \end{cases}$$




FIGURE 8. The picture of time- t slices when $|t| < 1/2$. They are the projections of the picture in Figure 7. As one can see, as $|t|$ increases, they would eventually cover themselves and the stalk will no longer be 1_V .

as in Example 3.4.11. We have considered the induced isotopy $\tilde{\Phi}$ on the universal cover $\dot{T}^*\mathbb{R}$ and its GKS kernel in Example 3.4.18. To get the time- t slices, one begins with the case when $|t| < 1/2$, in which case the same formula as in Example 3.4.18 applies. See Figure 8. For greater time t , one can use the decomposition $K(\Phi)|_t = K(\Phi)_{t/n}^{\circ n}$ to compute it. The claim is that a family version of the decomposition works, and that one can construct $K(\Phi)$ by first constructing its restriction $K(\Phi)|_{M \times M \times (-1/2, 1/2)}$ and extending it by convolution.

Before we leave this section, we mention that one main application of Theorem 3.4.15 is that it allows us to deform the underlying microsupport while keeping the category constant.

Proposition 3.4.42. *Let $\Phi : \dot{T}^*M \times I \rightarrow \dot{T}^*M$ be a homogeneous symplectic isotopy. Then, for any $F \in \text{Sh}(M)$ and $t \in I$, we have $\dot{\text{SS}}(K(\Phi)|_t \circ F) = \varphi_t(\dot{\text{SS}}(F))$. Thus, for any conic closed subset $X \subseteq \dot{T}^*M$, there is an equivalence*

$$\begin{aligned} \text{Sh}_{X \cup 0_M}(M) &\xrightarrow{\sim} \text{Sh}_{\varphi_t(X) \cup 0_M}(M) \\ F &\mapsto K(\Phi)|_t \circ F. \end{aligned}$$

Proof. A microsupport estimation similar to the one performed in the proof of Lemma 3.4.25 implies that $\dot{\text{SS}}(K(\Phi)|_t \circ F) \subseteq \varphi_t(\dot{\text{SS}}(F))$. But $K(\Phi)|_t$ has $K(\Phi^{-1})|_t$ as inverse so $\dot{\text{SS}}(F) = \dot{\text{SS}}(K(\Phi^{-1})|_t \circ K(\Phi)|_t \circ F) \subseteq \varphi_t^{-1}(\dot{\text{SS}}(K(\Phi)|_t \circ F))$ for the same reason. Applying φ_t , we obtain the reverse inclusion $\varphi_t(\dot{\text{SS}}(F)) \subseteq \dot{\text{SS}}(K(\Phi)|_t \circ F)$. This implies that $K(\Phi)|_t$ sends $\text{Sh}_{X \cup 0_M}(M)$ into $\text{Sh}_{\varphi_t(X) \cup 0_M}(M)$, but it is an equivalence, so we are done. $\square$

3.5. Microlocalization. The notion of support allows us to localize the category of sheaves on the base manifold M . Concretely speaking, for an open set $U \subseteq M$, the category $\text{Sh}(U)$ is the quotient of $\text{Sh}(M)$ by the subcategory $\text{Sh}(Z)$, where $Z := (M \setminus U)$ is the complementary subset. But we also know that $\text{Sh}(Z)$ has an alternative expression as $\text{Sh}_Z(M)$, the subcategory of sheaves on M consisting of those supported in Z . We could mimic this consideration and try to localize the category of sheaves on the cotangent bundle T^*M , or to simplify the wording, microlocalize the category by considering, for a conic open set $\Omega \subseteq \dot{T}^*M$, the quotient of $\text{Sh}(M)$ by $\text{Sh}_{\Omega^c}(M)$. Unlike the case on the base, this assignment is in general not a sheaf but only a presheaf. Its sheafification is usually denoted by μsh_{T^*M} when the context is clear, and a section of it is called a microsheaf.

Note that, as μsh_{T^*M} is conic, its restriction $\mu\text{sh}_{T^*M}|_{\dot{T}^*M}$ is necessarily a pullback $q^* \mu\text{sh}_{S^*M}$ for a unique sheaf μsh_{S^*M} along the quotient map $q : \dot{T}^*M \rightarrow S^*M$ to the cosphere bundle. The latter is the prototype for the local picture of *contact* geometry, an odd dimensional version of symplectic geometry. We will discuss that the notion of microsheaves in fact exists on general contact manifolds and this can in turn be used to define a notion of microsheaves on exact symplectic manifolds. We will begin the section by recalling some relevant details in order to construct the local category μsh_{T^*M} .

Lemma 3.5.1. *Let $\mathcal{C}$ be a stable category and $i : \mathcal{C}_0 \subseteq \mathcal{C}$ a stable subcategory. The quotient $p : \mathcal{C} \rightarrow \mathcal{C}/\mathcal{C}_0$ exists in st . Objects in $\mathcal{C}/\mathcal{C}_0$ are of the form $p(x)$ for $x \in \mathcal{C}$. Furthermore, for $x, y \in \mathcal{C}$, we have*

$$\text{Hom}_{\mathcal{C}/\mathcal{C}_0}(p(x), p(y)) = \underset{\substack{\alpha: y \rightarrow x^+ \\ \text{fib}(\alpha) \in \mathcal{C}_0}}{\text{colim}} \text{Hom}(x, y^+) = \underset{\substack{\beta: x^- \rightarrow x \\ \text{cof}(\beta) \in \mathcal{C}_0}}{\text{colim}} \text{Hom}(x^-, y)$$

where the first index runs over all morphisms $\alpha : y \rightarrow y^+$ such that the fiber $\text{fib}(\alpha)$ belongs to $\mathcal{C}_0$, and similarly for the second colimit.

Proof. We use the notation Pr_ω^R to denote the category of ω -presentable categories with morphisms given by limit-preserving functors that also preserve filtered colimits. By [44, Proposition 5.5.7.6], the inclusion $\text{Pr}_\omega^R \subseteq \widehat{\text{Cat}}$ preserves small limits. Thus, to compute the quotient $\mathcal{C}/\mathcal{C}_0$ in st , one can use Corollary 3.1.16, and compute $\text{Ind}(\mathcal{C})/\text{Ind}(\mathcal{C}_0)$ in $\text{Pr}_{\omega, \text{st}}^L$, which can again be computed as a limit in Pr_ω^R , which can be computed naively in $\widehat{\text{Cat}}$. More precisely, $\text{Ind}(i)$ has a right adjoint $\text{Ind}(i)^R$, and $\mathcal{C}/\mathcal{C}_0$ is given by the category of compact objects in $\text{fib}(\text{Ind}(i)^R)$. The description for the hom-anima follows from the fact that the Ind-completion $\text{Ind}(\mathcal{C})$ can be realized as $\text{Fun}^{\text{ex}}(\mathcal{C}^{\text{op}}, \mathcal{V})$ and the description comes from Kan extensions. $\square$

Remark 3.5.2. In the classical terminology, the category $\mathcal{C}/\mathcal{C}_0$ is the idempotent completion of the Verdier quotient.

Similar to the very large category of large categories $\widehat{\text{Cat}}$, we can consider the subcategory $\widehat{\text{st}} \subseteq \widehat{\text{Cat}}$ consisting of large stable categories. This category in particular contains the large category Pr_{st}^L as a non-full subcategory. By going up one universe higher, we can form quotients in st and have the same description for the hom-anima as in Lemma 3.5.1. For an inclusion $\mathcal{C}_0 \hookrightarrow \mathcal{C}$ in $\text{Pr}_{\text{st}}^L \cap \text{Pr}_{\text{st}}^R$, we thus have three ways to form the quotient $\mathcal{C}_1 := \mathcal{C}/\mathcal{C}_0$, and the next two lemmas assert that they are all the same. Recall that we defined the notion of recollement in Definition 3.1.25.

Lemma 3.5.3. *Let $\iota_* : \mathcal{C}_0 \hookrightarrow \mathcal{C}$ be an inclusion in $\text{Pr}_{\text{st}}^L \cap \text{Pr}_{\text{st}}^R$. That is, ι_* admits both a left adjoint ι^* and a right adjoint $\iota^!$. Then, the quotient $\mathcal{C}_1 := \mathcal{C}/\mathcal{C}_0$ taken in Pr_{st}^L , Pr_{st}^R , and $\widehat{\text{st}}$ all agree, and $(\mathcal{C}_0, \mathcal{C}_1)$ forms a recollement of $\mathcal{C}$.*

Proof. The quotient in Pr_{st}^L and Pr_{st}^R can both be computed as limits in $\widehat{\text{Cat}}$ by passing to the corresponding adjoints, via the fact that $\text{Pr}_{\text{st}}^L = (\text{Pr}_{\text{st}}^R)^{\text{op}}$ and the fact that both $\text{Pr}_{\text{st}}^L, \text{Pr}_{\text{st}}^R \subseteq \widehat{\text{Cat}}$ preserve limits, by Proposition 2.2.24. This would naturally imply that $\mathcal{C}_1$ can be viewed both as the left and the right orthogonal complement via the corresponding adjoints of the quotient $j^* : \mathcal{C} \rightarrow \mathcal{C}_1$ and the general structure of recollement will imply that they are equivalent. To compare them with the quotient in $\widehat{\text{st}}$, temporarily denoted by $\mathcal{C}'_1$, we note that the universal property implies that there is a canonical map $\mathcal{C}'_1 \rightarrow \mathcal{C}_1$ as

$j^* \iota_* = 0$, that is clearly surjective. For $x, y \in \mathcal{C}$, Lemma 3.5.1 implies that $\mathrm{Hom}_{\mathcal{C}_1}(p(x), p(y))$ is computed by $\mathrm{Hom}_{\mathcal{C}}(x, y^+)$ running over $\alpha : y \rightarrow y^+$ such that $\mathrm{fib}(\alpha) \in \mathcal{C}_0$. But the fiber sequence, $\iota_* \iota^! y \rightarrow y \xrightarrow{\eta_y} j_* j^* y$ implies η_y is cofinal among such morphisms so

$$\mathrm{Hom}_{\mathcal{C}'_1}(p(x), p(y)) = \mathrm{colim}_{\substack{\alpha: y \rightarrow y' \\ \mathrm{fib}(\alpha) \in \mathcal{C}_0}} \mathrm{Hom}_{\mathcal{C}}(x, y^+) \xrightarrow{\sim} \mathrm{Hom}_{\mathcal{C}}(x, j_* j^* y) = \mathrm{Hom}_{\mathcal{C}_1}(j^* x, j^* y).$$

□

Before we define the notion of microsheaves, we mention that sheafification exists for $\widehat{\mathrm{st}}$ -valued sheaves.

Lemma 3.5.4. *Let X be a topological space. The sheafification functor $(-)^{\dagger}$, the left adjoint of the inclusion $\mathrm{Sh}(X; \mathrm{st}) \hookrightarrow \mathrm{PSh}(X; \mathrm{st})$, exists. Similarly, the sheafification functor exists over $\widehat{\mathrm{st}}$.*

Proof. By going one universe higher, it suffices to prove the case for st . By [7, Corollary 4.25], st is presentable, in fact, ω -presentable. By [33, Proposition 4.7], this cited corollary also implies the version over a general coefficient $\mathcal{V} \in \mathrm{CAlg}(\mathrm{Pr}_{\omega, \mathrm{st}}^{\mathrm{L}})$. Alternatively, one could follow the proof strategy presented in [56, Theorem 2.9]. Then, by (2) of Proposition 2.5.5 and Lemma 2.5.11, $\mathrm{Sh}(X; \mathrm{st}) = \mathrm{Fun}^R(\mathrm{Sh}(X; \mathrm{Ani})^{\mathrm{op}}, \mathrm{st}) = \mathrm{Sh}(X; \mathrm{Ani}) \otimes \mathrm{st}$ is presentable and similarly for $\mathrm{PSh}(X; \mathrm{st})$. Tracing through the identification, we see that $(-)^{\dagger, \mathrm{Ani}} \otimes \mathrm{id}_{\mathrm{st}}$ corresponds to the sheafification functor $(-)^{\dagger, \mathrm{st}}$ over st . □

Definition 3.5.5. Let M be a manifold. We define the $\widehat{\mathrm{st}}$ -valued presheaf $\mu\mathrm{sh}_{T^*M}^{\mathrm{pre}}$ by

$$\begin{aligned} \mu\mathrm{sh}_{T^*M}^{\mathrm{pre}} : \mathrm{Op}_{T^*M}^{\mathrm{op}} &\rightarrow \widehat{\mathrm{st}} \\ \Omega &\mapsto \mathrm{Sh}(M)/\mathrm{Sh}_{\Omega^c}(M). \end{aligned}$$

Note that, concretely, $\mathrm{Sh}_{\Omega^c}(M) = \{F \mid \mathrm{SS}(F) \cap \Omega = \emptyset\}$ is given by sheaves microsupported outside Ω . We then define $\mu\mathrm{sh}_{T^*M}$ to be the sheafification of $\mu\mathrm{sh}_{T^*M}^{\mathrm{pre}}$ over $\widehat{\mathrm{st}}$. As $\mu\mathrm{sh}_{T^*M}$ is conic, since $\mathrm{SS}(F)$ is conic for any $F \in \mathrm{Sh}(M)$, the restriction away from the zero section, $\mu\mathrm{sh}_{T^*M}|_{\dot{T}^*M} = q^* \mu\mathrm{sh}_{S^*M}$, is necessarily a pullback of a unique sheaf $\mu\mathrm{sh}_{S^*M}$ on S^*M , which could also be defined using a similar presheaf $\mu\mathrm{sh}_{S^*M}^{\mathrm{pre}}$. For a conic open set $\Omega \subseteq T^*M$, we call an object $\mathcal{F} \in \mu\mathrm{sh}_{T^*M}(\Omega)$ a “microsheaf” on Ω so we sometimes refer to $\mu\mathrm{sh}_{T^*M}$ as the (category-valued) sheaf of microsheaves. We use the same notation for sections of $\mu\mathrm{sh}_{S^*M}$. When the context is clear, we sometimes drop the subscript.

Lemma 3.5.6. *Denote by $\pi : T^*M \rightarrow M$ the projection. Then, $\pi_* \mu\mathrm{sh} = \mathrm{Sh}^*$, the assignment $(U \mapsto \mathrm{Sh}(U))$ with transition map given by $*$ -pullback.*

Proof. Let $s_0 : M \hookrightarrow T^*M$ be the zero section. Since $\mu\mathrm{sh}$ is conic, we have $\pi_* \mu\mathrm{sh} = s_0^* \mu\mathrm{sh}$ by an argument similar to [36, Proposition 3.7.2]. But s_0^* is a left adjoint so $s_0^* \mu\mathrm{sh} = (s_0^{*, \mathrm{pre}} \mu\mathrm{sh}^{\mathrm{pre}})^{\dagger}$. Again, being conic implies that $s_0^{*, \mathrm{pre}} \mu\mathrm{sh}^{\mathrm{pre}} = \pi_*^{\mathrm{pre}} \mu\mathrm{sh}^{\mathrm{pre}}$. Now recall that, if we set $Z := (M \setminus U)$, then Example 3.2.11 and Proposition 3.1.23 imply that $\mathrm{Sh}_{T^*M|Z}(M) = \mathrm{Sh}(Z)$. We then observe that, by Lemma 3.1.28,

$$\pi_*^{\mathrm{pre}} \mu\mathrm{sh}^{\mathrm{pre}}(U) = \mu\mathrm{sh}^{\mathrm{pre}}(T^*U) = \mathrm{Sh}(M)/\mathrm{Sh}_{T^*M|Z}(M) = \mathrm{Sh}(M)/\mathrm{Sh}(Z) = \mathrm{Sh}(U)$$

which we know is already a sheaf by Proposition 2.3.18. □

Outside of the constructible situation, which we will see later in Proposition 3.5.27, little is known about a general microsheaf $\mathcal{F}$. It is however slightly more manageable at stalks.

Lemma 3.5.7. *Let $(x, \xi) \in T^*M$ be a covector. Then, objects in the stalk $\mu\text{sh}_{(x, \xi)}$ are given by pairs (Ω, F_Ω) such that $F_\Omega \in \mu\text{sh}^{\text{pre}}(\Omega)$ where $\Omega \subseteq T^*M$ is a conic subset. If $(\Omega', G_{\Omega'})$ is another pair, then*

$$\text{Hom}_{\mu\text{sh}_{(x, \xi)}}((\Omega', G_{\Omega'}), (\Omega, F_\Omega)) = \text{colim}_{\Omega'' \ni (x, \xi)} \text{Hom}_{\mu\text{sh}^{\text{pre}}(\Omega'')} (G_{\Omega'}|_{\Omega''}, F_\Omega|_{\Omega''})$$

where Ω'' runs over conic open sets contained in $\Omega \cap \Omega'$ and we use the notation $F_\Omega|_{\Omega''}$ to denote its image under the projection $\mu\text{sh}^{\text{pre}}(\Omega) \rightarrow \mu\text{sh}^{\text{pre}}(\Omega'')$.

Proof. The main observation is that, since sheafification $(-)^{\dagger}$ and taking stalks $i_{(x, \xi)}^*$ are both right adjoints, one can commute them and get $i_{(x, \xi)}^{*, \text{pre}}$, which means that the stalk can be computed at the level of presheaves. The rest is then Proposition 2.1.12. $\square$

Example 3.5.8. Note that for $x \in M$, the functor $\text{Sh}(M) \rightarrow (\text{Sh}^*)_x$ to the stalk of Sh^* is different from taking the stalk $i_x^* : \text{Sh}(M) \rightarrow \mathcal{V}$. Indeed, objects in $(\text{Sh}^*)_x$ are germs of sheaves near x , so there is an induced functor $i_x^* : (\text{Sh}^*)_x \rightarrow \mathcal{V}$ sending a pair (U, F) , $F \in \text{Sh}(U)$, to the stalk F_x . This map is not an isomorphism in general. For example, take $X = \mathbb{R}^1$. By Lemma 3.3.13, $\text{Hom}(1_{\{0\}}, 1_{\mathbb{R}^1}) = 1_{\mathcal{V}}[-1]$. As intervals of the form $(-\epsilon, \epsilon)$ are cofinal, the $\mathcal{V}$ -valued hom in $(\text{Sh}^*)_x$ is also given by $1_{\mathcal{V}}$. However, $i_0^* 1_{\{0\}} = i_0^* 1_{\mathbb{R}^1} = 1_{\mathcal{V}}$, so the $\mathcal{V}$ -valued hom between them (computed at the level of stalks) is $1_{\mathcal{V}}$. Said differently, taking stalks does not commute with $\mathcal{H}\text{om}$ in general.

Lemma 3.5.9. *If $\mathcal{C}$ is a sheaf of categories on X , then for sections $s_1, s_2 \in \mathcal{C}(U)$ on U , the assignment*

$$\begin{aligned} \text{Op}_U^{\text{op}} &\rightarrow \text{Ani} \\ V &\mapsto \text{Hom}_{\mathcal{C}(V)}(s_1|_V, s_2|_V) \end{aligned}$$

forms a sheaf, which we will denote by $\mathcal{H}\text{om}(s_1, s_2) \in \text{Sh}(X; \text{Ani})$. The same statement holds with coefficients.

Proof. The statement over anima is a direct consequence of Remark 2.1.11 from which one can deduce the general statement. $\square$

Definition 3.5.10. Let $\mathcal{C}$ be a sheaf valued in a $\mathcal{V}$ -enriched category on X . For $s \in \mathcal{C}(U)$, the support of s , $\text{supp}(s)$, is defined to be the complement of the largest open set such that $s|_U = 0$.

Lemma 3.5.11. *Let s be as in Definition 3.5.10. Then, we have $\text{supp}(s) = \text{supp}(\mathcal{H}\text{om}(s, s))$.*

Proof. This follows from the fact that, in a stable category, an object x is zero if and only if the identity id_x is zero. $\square$

The next thing we will discuss is that the support of microsheaves is related to microsupport.

Lemma 3.5.12. *There is a well-defined notion of microsupport on Ω for objects in $\mu\text{sh}^{\text{pre}}(\Omega)$. More precisely, let $\Omega \subseteq T^*M$ be a conic open set, and we use the notation*

$$\text{SS}_{\Omega}(F) := \text{SS}(F) \cap \Omega.$$

Denote by $q^ : \text{Sh}(M) \rightarrow \mu\text{sh}^{\text{pre}}(\Omega) = \text{Sh}(M)/\text{Sh}_{\Omega^c}(M)$ the quotient map. If $F_1, F_2 \in \text{Sh}(M)$ are such that $q^*F_1 = q^*F_2$, then we have $\text{SS}_{\Omega}(F_1) = \text{SS}_{\Omega}(F_2)$. As a special case, $q^*(F) = 0$ in $\mu\text{sh}^{\text{pre}}(\Omega)$ if and only if $\text{SS}_{\Omega}(F) = \emptyset$.*

Proof. By Lemma 3.5.3, $(\mathrm{Sh}_{\Omega^c}(M), \mu\mathrm{sh}^{\mathrm{pre}}(\Omega))$ forms a recollement of $\mathrm{Sh}(M)$. Thus, for $F \in \mathrm{Sh}(M)$, we have the fiber sequence $q_!q^*F \rightarrow F \rightarrow \iota_*\iota^*F$. By definition, $\mathrm{SS}_{\Omega}(\iota_*\iota^*F) = \emptyset$. The microlocal triangular inequality, Proposition 3.3.2, thus implies that $\mathrm{SS}_{\Omega}(F) = \mathrm{SS}_{\Omega}(q_!q^*F)$. But then, we have $\mathrm{SS}_{\Omega}(F_1) = \mathrm{SS}_{\Omega}(q_!q^*F_1) = \mathrm{SS}_{\Omega}(q_!q^*F_2) = \mathrm{SS}_{\Omega}(F_2)$ as desired. $\square$

The following proposition is a direct consequence of the above Lemma 3.5.12 and tracing through definitions.

Proposition 3.5.13. *Let $(x, \xi) \in T^*M$ be a point and $\Omega \subseteq T^*M$ a conic open set. For $\mathcal{F} \in \mu\mathrm{sh}_{T^*M}(\Omega)$, we have $(x, \xi) \notin \mathrm{supp}(\mathcal{F})$ if and only if there exists (and hence for all) a representative $(\Omega_0, F_{\Omega_0}) \in (\mu\mathrm{sh}_{T^*M})_{(x, \xi)}$ such that $(x, \xi) \notin \mathrm{SS}_{\Omega_0}(F_{\Omega_0})$.*

Oftentimes, the sheaf $\mu\mathrm{sh}_{T^*M}$ is too large to work with, and one restricts to a smaller sheaf by prescribing the support condition. In what follows, we will consider, in a sense, the class of smallest possible subsheaves of $\mu\mathrm{sh}_{T^*M}$.

Definition 3.5.14. Let (V, ω) be a finite dimensional symplectic vector space, i.e., ω is a non-degenerate alternating bilinear form.

- (1) Let $W \subseteq V$ be a subvector space. The symplectic complement (with respect to ω) of W is the subvector space

$$W^{\omega} := \{v \in V \mid \omega(v, w) = 0 \text{ for all } w \in W\}.$$

- (2) A subspace $C \subseteq V$ is said to be coisotropic if $C^{\omega} \subseteq C$.
- (3) A subspace $I \subseteq V$ is said to be isotropic if $I \subseteq I^{\omega}$.
- (4) A subspace $L \subseteq V$ is said to be Lagrangian if L is both isotropic and coisotropic.

Remark 3.5.15. In a finite dimensional symplectic vector space (V, ω) , one can check that

- (1) $(W^{\omega})^{\omega} = W$,
- (2) a coisotropic subspace C satisfies $\dim C \geq \frac{1}{2} \dim V$,
- (3) an isotropic subspace I satisfies $\dim I \leq \frac{1}{2} \dim V$,
- (4) a Lagrangian subspace L satisfies $\dim L = \frac{1}{2} \dim V$.

Example 3.5.16. Consider $\mathbb{R}^4$ with basis given by $\{e_1, f_1, e_2, f_2\}$ and $\omega = e_1^* \wedge f_1^* + e_2^* \wedge f_2^*$. In this case, $C := \langle e_1, f_1, e_2 \rangle$ has symplectic complement $C^{\omega} = \langle e_2 \rangle$, and is thus coisotropic. Similarly, one can check that $L = \langle e_1, e_2 \rangle$ is Lagrangian.

Definition 3.5.17. Let (X, ω) be a symplectic manifold. We say that a submanifold $C \subseteq X$ is coisotropic if for every point $x \in X$, the subspace $T_x C \subseteq T_x X$ is a coisotropic subvector space. We say that a submanifold $L \subseteq X$ is Lagrangian if L is coisotropic and $\dim L = \frac{1}{2} \dim X$.

Example 3.5.18. Let M be a smooth manifold. Recall that, as we saw in Example 3.4.9, the Liouville form λ on T^*M , given locally on a chart (x, ξ) by $\lambda = \sum_{k=1}^n \xi_k dx_k$, induces a symplectic form $\omega = d\lambda$. Two standard examples of Lagrangians are the zero section 0_M and the cotangent fibers T_x^*M .

Definition 3.5.19. Assume the underlying manifold X of a symplectic manifold (X, ω) is C^{ω} , i.e., real analytic. We say a subanalytic set $\Lambda \subseteq X$ is a (singular) Lagrangian, a Lagrangian variety, or simply Lagrangian, if Λ has a Lagrangian dense open stratum. More precisely, there exists a Whitney stratification $\Lambda = \cup_{\alpha \in \mathcal{S}} L_{\alpha}$ by locally closed real analytic submanifolds L_{α} such that the top dimensional strata, which are necessarily open in Λ , are Lagrangian submanifolds.

Example 3.5.20. In $T^*\mathbb{R}^1$ with coordinate (x, ξ) , the set $\{x\xi = 0\} = 0_{\mathbb{R}^1} \cup T_0^*\mathbb{R}^1$ is Lagrangian. Indeed, it can be stratified by

$$\{(-\infty, 0), (0, \infty), T_{0, >}^*\mathbb{R}^1, T_{0, <}^*\mathbb{R}^1, \{0\}\}.$$

In fact, if $M = \coprod_{\alpha \in \mathcal{S}} X_\alpha$ is a real analytic manifold with a Whitney stratification $\mathcal{S}$, then its conormal bundle

$$(16) \quad N^*(\mathcal{S}) := \coprod_{\alpha \in \mathcal{S}} N^*(X_\alpha)$$

is Lagrangian.

Theorem 3.5.21 ([36, Theorem 6.5.4, Proposition 8.3.10]).

- (1) *Let M be a manifold and $F \in \text{Sh}(M)$ be a sheaf. The microsupport $\text{SS}(F) \subseteq T^*M$ is coisotropic.*
- (2) *Assume further that M is real analytic and $\text{SS}(F)$ is subanalytic. Then, F is constructible, in the sense of Definition 3.4.30, if and only if $\text{SS}(F)$ is Lagrangian.*

The above theorem is a consequence of the following propositions:

Proposition 3.5.22 ([36, Corollary 8.3.22]). *Let M be real analytic and $\Lambda \subseteq T^*M$ be a conic subanalytic Lagrangian. Then, there exists Whitney stratification $\mathcal{S}$ such that $\Lambda \subseteq N^*(\mathcal{S})$.*

Proposition 3.5.23 ([26, proposition 4.8]). *Let $\mathcal{S}$ be a Whitney stratification and denote by $\text{Sh}_{\mathcal{S}}(M)$ the subcategory of $\mathcal{S}$ -constructible sheaves. Then we have $\text{Sh}_{\mathcal{S}}(M) = \text{Sh}_{N^*(\mathcal{S})}(M)$.*

Corollary 3.5.24. *Let M be real analytic and $\Lambda \subseteq T^*M$ be a conic subanalytic Lagrangian. Then, the category $\text{Sh}_{\Lambda}(M)$ is ω -accessible.*

Back to the discussion of microsheaves.

Definition 3.5.25. Let $C \subseteq T^*M$ be a closed subset. We use the notation $\mu\text{sh}_{T^*M, C}$ or simply μsh_C to denote the subsheaf of μsh_{T^*M} consisting of microsheaves $\mathcal{F}$ such that $\text{supp}(\mathcal{F}) \subseteq C$.

In the case when $C = \Lambda$ is a subanalytic Lagrangian. There is an alternative way to obtain the same sheaf.

Notation 3.5.26. We denote by

$$\begin{aligned} \mu\text{sh}_{\Lambda}^{\text{pre}} : \text{Op}_{T^*M}^{\text{op}} &\rightarrow \widehat{\text{st}} \\ \Omega &\mapsto \text{Sh}_{\Lambda \cup \Omega^c}(M) / \text{Sh}_{\Omega^c}(M). \end{aligned}$$

Writing out the definition, $\mu\text{sh}_{\Lambda}^{\text{pre}}(\Omega)$ is given by $\{F \mid \text{SS}_{\Omega}(F) \subseteq \Lambda\}$, those sheaves microsupported in Λ on Ω , quotiented by $\{F \mid \text{SS}_{\Omega}(F) = \emptyset\}$.

Proposition 3.5.27. *Let M be real analytic and $\Lambda \subseteq T^*M$ be a conic subanalytic Lagrangian. Then,*

- (1) ([13, 3.1.1]) *The canonical map $\mu\text{sh}_{\Lambda}^{\text{pre}} \rightarrow \mu\text{sh}_{\Lambda}$ induces an equivalence*

$$(\mu\text{sh}_{\Lambda}^{\text{pre}})^{\dagger} \xrightarrow{\sim} \mu\text{sh}_{\Lambda}.$$

- (2) *$\mu\text{sh}_{\Lambda}^{\text{pre}}$ is constructible. In particular, for any $(x, \xi) \in T^*M$ and small enough $\Omega \subseteq T^*M$,*

$$\mu\text{sh}_{\Lambda}(\Omega) = \mu\text{sh}_{\Lambda}^{\text{pre}}(\Omega) = \text{Sh}_{\Lambda}(B) / \text{Sh}_{\Lambda \cap \Omega^c}(M)$$

where $B = \pi(\Omega)$ is the projection of Ω to the zero section.

- (3) μsh_Λ is a sheaf valued in $\text{Pr}_{\omega, st}^R$, i.e., μsh_Λ takes values in ω -accessible categories, and the restriction maps are both limit-preserving and colimit-preserving.
- (4) For any $\mathcal{F} \in \mu\text{sh}_\Lambda(\Omega)$ and $(x, \xi) \in \Omega$, there exists a small Ω_0 and $F \in \mu\text{sh}_\Lambda^{\text{pre}}(\Omega_0)$ such that $\text{supp}(\mathcal{F}) \cap \Omega_0 = \text{SS}_{\Omega_0}(F)$.

3.6. Globalizing microsheaves. Our next goal is to extend the notion of microsheaves to an arbitrary co-orientable contact manifold. Let V be a C^∞ -manifold of dimension $2n + 1$, $n \geq 0$, and $\mathcal{H}$ be a hyperplane distribution, i.e., a subbundle $\mathcal{H} \subseteq TV$ of rank $2n$. Consider the short exact sequence of vector bundles over V , $0 \rightarrow \mathcal{H} \rightarrow TV \rightarrow TV/\mathcal{H} \rightarrow 0$, and the dual sequence

$$0 \rightarrow (TV/\mathcal{H})^\vee \xrightarrow{i} T^*V \rightarrow \mathcal{H}^\vee \rightarrow 0.$$

Definition 3.6.1. The pair $(V, \mathcal{H})$ is said to be a contact manifold (with contact distribution $\mathcal{H}$) if $i : (TV/\mathcal{H})^\vee \subseteq T^*V$ is a symplectic submanifold. That is, the symplectic form ω from Example 3.4.9 restricts to a symplectic form $i^*\omega$ on $(TV/\mathcal{H})^\vee$. In this case, we denote $\tilde{V} := (TV/\mathcal{H})^\vee$ and call it the symplectization of V . We sometimes abuse notation and simply say that V is a contact manifold without explicitly naming its distribution $\mathcal{H}$.

Remark 3.6.2. Alternatively, since $\mathcal{H}$ is a hyperplane distribution, it is locally of the form $\ker \alpha$ for some 1-form α . Then, $(V, \mathcal{H})$ is a contact manifold if and only if $\alpha \wedge (d\alpha)^n$ is non-vanishing near any point, for any local form α . Equivalently, $d\alpha$ restricts to a symplectic form on $\mathcal{H}$ for such a local form.

Definition 3.6.3. A contact manifold $(V, \mathcal{H})$ is co-oriented by α if $\mathcal{H} = \ker \alpha$ for some global 1-form $\alpha \in \Omega^1(V)$.

Example 3.6.4. The cosphere bundle S^*M is a co-orientable contact manifold with its contact distribution $\mathcal{H}$ given by $\mathcal{H}_{[x, \xi]} = \ker(\sum_{k=1}^n \xi_k dx_k)$. To co-orient it, one can choose a Riemannian metric g on M and embed S^*M into $\dot{T}^*M$ by a section $s : S^*M \hookrightarrow \dot{T}^*M$ as norm-1 covectors and pullback the standard Liouville form. Similarly, the coprojectivization $\mathbb{P}^*M := \dot{T}^*M/\mathbb{R}^\times$, which is a quotient of S^*M by the $\mathbb{Z}/2\mathbb{Z}$ -action given by the antipodal map, is a contact manifold, but it is not always co-orientable.

Definition 3.6.5. A C^∞ map $f : (V, \mathcal{H}) \hookrightarrow (V', \mathcal{H}')$ is said to be a contact embedding if $f(\mathcal{H}) = f(TV) \cap \mathcal{H}'$. Alternatively, this means that if $\mathcal{H}' = \ker \alpha'$ for some 1-form α' , then $\ker(f^*\alpha) = \mathcal{H}$.

Example 3.6.6. The (1-)jet bundle J^1M of a manifold M is the product manifold $T^*M \times \mathbb{R}$. Call the coordinate of $\mathbb{R}$ by s , then the 1-form $\alpha = \lambda + ds$ is a contact form where λ is the standard Liouville form on T^*M . The map

$$(17) \quad \begin{aligned} J^1M &\hookrightarrow S^*(M \times \mathbb{R}) \\ (x, \xi, s) &\mapsto (x, s, [\xi, 1]) \end{aligned}$$

is an open contact embedding.

In fact, contact geometry is “locally trivial” in the sense that any contact manifold is locally contactomorphic to the cosphere bundle from Example 3.6.4.

Proposition 3.6.7. Let V be a manifold. Then, for any point $p \in V$, there exists a neighborhood $\mathcal{U}$ of p and an open contact embedding $f : \mathcal{U} \hookrightarrow S^*M$ for some M such that $\dim M = \frac{1}{2}(\dim V + 1)$. Such a pair $(\mathcal{U}, f)$ is called a Darboux chart of V .

Our goal is to define sheaves of categories on V that are equivalent to μsh_{S^*M} locally on Darboux charts. The reason we switch from symplectic geometry to contact geometry is that the latter satisfies Gromov's h-principle [19].

Theorem 3.6.8 (Gromov's h-principle). *Let V be a co-orientable contact manifold. Then, there exists $N \in \mathbb{N}$ and a contact embedding $f : V \hookrightarrow J^1\mathbb{R}^N$. Furthermore, the space of contact embeddings can be made arbitrarily connected by stabilizing, that is, by further embedding into $J^1\mathbb{R}^{N+k}$ along the stabilization map, viewing $J^1\mathbb{R}^{N+k} \cong J^1\mathbb{R}^N \times T^*\mathbb{R}^k$,*

$$(18) \quad \begin{aligned} J^1\mathbb{R}^N &\hookrightarrow J^1\mathbb{R}^N \times T^*\mathbb{R}^k \\ (x, \xi, s) &\mapsto (x, \xi, s, 0, 0), \end{aligned}$$

for $k \gg 1$.

Example 3.6.9. Let M be a C^∞ manifold and choose a smooth embedding $g : M \hookrightarrow \mathbb{R}^N$ for some $N \gg 1$. This embedding induces a short exact sequence

$$0 \rightarrow N^*(M) \rightarrow T^*\mathbb{R}^N|_M \xrightarrow{dg^\vee} T^*M \rightarrow 0$$

of vector bundles over M , which always splits, i.e., there is a section $f : T^*M \hookrightarrow T^*\mathbb{R}^N|_M$. One can check that the composition $T^*M \xrightarrow{f} T^*\mathbb{R}^N|_M \hookrightarrow T^*\mathbb{R}^N$ is a homogeneous symplectic embedding and lands in $T^*\mathbb{R}^N$. Thus, it descends to a contact embedding $S^*M \hookrightarrow S^*\mathbb{R}^N$. We remark that the space of sections is convex. That is, if f_0, f_1 are both sections of $dg^\vee$, then $f_t := tf_0 + (t-1)f_1$ is also a section for $t \in [0, 1]$.

Fix once and for all an open embedding of $J^1M \hookrightarrow S^*(M \times \mathbb{R})$ that is compatible with the stabilization map (18), for example, the one considered in Example 3.6.6. One naive definition of “microsheaves on V ” might be $\mu\text{sh}_{S^*(\mathbb{R}^{N+1}), \text{im}f}$ for such a large-codimension embedding f given by Gromov's h-principle. However, Theorem 3.5.21 will imply that the sheaf is most likely zero if $N \gg 1$, and we will need N to be large anyway, since the connectivity condition of the embedding space is needed to show the independence of the candidate sheaf. The solution is to consider a notion of “thickening”, which we will first introduce with an example before giving the general definition.

Example 3.6.10. Consider the simplest non-trivial example of Example 3.6.9 given by $i : \mathbb{R}^1 \hookrightarrow \mathbb{R}^3$ as the x -axis. Let $(x, \xi, y, \eta, z, \zeta)$ be the coordinate of $T^*\mathbb{R}^3$. A section of the differential on the cotangent bundle is given by

$$\begin{aligned} f : T^*\mathbb{R}^1 &\hookrightarrow T^*\mathbb{R}^3 \\ (x, \xi) &\mapsto (x, \xi, 0, 0, 0, 0). \end{aligned}$$

We will consider the thickening issue on the level of sheaves, but the discussion regarding microsheaves can be made similarly.

First, we notice that the category $\text{Sh}_{\text{im}f}(\mathbb{R}^3)$ is 0: If $\text{SS}(F) \subseteq \text{im}(f)$, it in particular implies that $\text{supp}(F) \subseteq \mathbb{R}^1$, so by Proposition 3.1.23, $F = i_*G$ for some $G \in \text{Sh}(\mathbb{R}^1)$. Since i is proper, by Proposition 3.3.14, the microsupport estimation $\text{SS}(i_*G) = i_\pi((di^\vee)^{-1}\text{SS}(G))$ so if there is any point $(x, \xi) \in \text{SS}(G)$, we would have $(x, \xi, 0, \eta, 0, \zeta) \in \text{SS}(F)$ for all $\eta, \zeta \in \mathbb{R}$, which we are assuming is not the case. Thus, $\text{SS}(G) = \emptyset$ and $G = 0$.

Now, we explain the thickening. The normal bundle of the embedding f , $N(f) \cong T^*\mathbb{R}^2$ is a symplectic vector bundle with the standard symplectic structure $d\eta \wedge dy + d\zeta \wedge dz$. There are two natural choices of Lagrangian subbundles, the first is spanned by $\{\frac{\partial}{\partial y}, \frac{\partial}{\partial z}\}$, and the other

is spanned by $\{\frac{\partial}{\partial\eta}, \frac{\partial}{\partial\zeta}\}$. The first comes from the coisotropic subspace $N^*(\mathbb{R}^1) \subseteq T^*\mathbb{R}^3$, given by the conormal bundle of $\mathbb{R}^1$ inside $\mathbb{R}^3$. More concretely, it is given by $\{(x, \xi, 0, \eta, 0, \zeta)\}$. By Example 3.2.11, we have that the map

$$\begin{aligned} i_* \text{Sh}(\mathbb{R}^1) &\rightarrow \text{Sh}_{N^*(\mathbb{R}^1)}(\mathbb{R}^3) \\ F &\mapsto i_* F \end{aligned}$$

is an equivalence so the thickening $N^*(\mathbb{R}^1)$ recovers $\text{Sh}(\mathbb{R}^1)$. The second Lagrangian subbundle gives rise to the coisotropic subspace $T^*\mathbb{R}^1 \times 0_{\mathbb{R}^2} = \{(x, \xi, y, 0, z, \zeta)\}$. Again, we recover $\text{Sh}(\mathbb{R}^1)$ using Lemma 3.4.28, which states that the map

$$\begin{aligned} p^* \text{Sh}(\mathbb{R}^1) &\rightarrow \text{Sh}_{T^*\mathbb{R}^1 \times 0_{\mathbb{R}^2}}(\mathbb{R}^3) \\ F &\mapsto p^* F \end{aligned}$$

where $p : \mathbb{R}^3 \rightarrow \mathbb{R}^1$ is the projection given by $p(x, y, z) = x$, is an equivalence.

Now, we recall if $C \subseteq (X, \omega)$ is a coisotropic submanifold, then one can check that the distribution $(TC)^\omega$, which is in fact a subbundle, given at $x \in X$ by $(T_x C)^\omega \subseteq TX$, the symplectic complement of $T_x C$, is involutive. That is, if X, Y are vector fields in $(TC)^\omega$, then their Lie bracket $[X, Y]$ will still be in $(TC)^\omega$. By Frobenius theorem, the manifold C is foliated by isotropic submanifolds of dimension $\dim X - \dim C$; we usually refer to it as the characteristic foliation associated to C and denote it by $\mathcal{F}_C$. A similar discussion holds in the contact setting.

Definition 3.6.11. Let $f : V \subseteq W$ be a contact embedding with $\dim V = 2n + 1$ and $\dim W = 2(n+k)+1$. A thickening V^σ of V is a coisotropic submanifold $V^\sigma \subseteq W$ containing V , such that $\dim V^\sigma = 2n+k+1$, and such that the characteristic foliation $\mathcal{F}_{V^\sigma}$ is transversal to V .

The reason to use the notation V^σ is the following: Since f is a contact embedding, its normal bundle $N(f)$ is naturally a symplectic vector bundle over V . As the characteristic foliation is transversal to V , it defines an isotropic subbundle $\sigma \subseteq N(f)$ and, by the dimension assumption, σ is in fact a Lagrangian subbundle. In other words, there is a map

$$(19) \quad \begin{aligned} \mathfrak{Thick}(V \subseteq W) &\rightarrow \mathfrak{Lag}(N(f)) \\ V^\sigma &\mapsto \sigma \end{aligned}$$

from the space of thickenings to the space of Lagrangian subbundles of $N(f)$, and it is in fact a homotopy equivalence. See [12, APPENDIX C] for more details. Concretely, for example, a homotopy $\{\sigma_t\}_{t \in [0,1]}$ between two Lagrangian subbundles σ_1 and σ_2 can be realized by a family of coisotropic submanifolds V^{σ_t} .

With the notion of thickening, one begins with a large-codimension embedding $f : V \hookrightarrow J^1\mathbb{R}^N (\subseteq S^*(\mathbb{R}^{N+1}))$ and can consider the subsheaf $\mu\text{sh}_{V^\sigma} \equiv \mu\text{sh}_{S^*(\mathbb{R}^{N+1}), V^\sigma} \subseteq \mu\text{sh}_{S^*(\mathbb{R}^{N+1})}$ consisting of microsheaves supported in V^σ and then, since we want to have a sheaf on V , consider its restriction $\mu\text{sh}_{V^\sigma}|_V$. Our first task is to consider the dependence of this sheaf on the embedding f and the thickening:

Lemma 3.6.12. *Let V be a co-oriented contact manifold and choose an embedding $f : V \hookrightarrow J^1\mathbb{R}^N$ for some $N \gg 1$. The sheaf of categories $\mu\text{sh}_{V, -\sigma} := \mu\text{sh}_{V^\sigma}|_V$ depends only on the*

choice of the (stable) homotopy class $[\sigma] \subseteq N(f)$ of a Lagrangian subbundle of the normal bundle $N(f)$.³²

Sketch of proof. The first fact we use concerns the stabilization map (18). More precisely, when further embedding $f : V \subseteq J^1\mathbb{R}^N$ into $J^1\mathbb{R}^{N+k}$ via (18), one can consider the thickening $V^\sigma \times T_0^*\mathbb{R}^k$ and the claim is that the map $i_* : \text{Sh}(\mathbb{R}^{N+1}) \hookrightarrow \text{Sh}(\mathbb{R}^{N+k+1})$, where $i : \mathbb{R}^{N+1} \rightarrow \mathbb{R}^{N+k+1}$ is given by $(x, s) \mapsto (x, y, s)$, induces an equivalence

$$[(di)^\vee]^* \mu\text{sh}_{T^*\mathbb{R}^{N+1}} = \mu\text{sh}_{T^*\mathbb{R}^{N+k+1}}|_{\mathbb{R}^{N+1}},$$

see for example [11, Lemma 4.9], so the subsheaf $\mu\text{sh}_{V^\sigma \times T_0^*\mathbb{R}^k}|_V$ remains unchanged.

Then, the second fact we will use is that the sheaf of microsheaves μsh satisfies a constancy property similar to the one discussed in Proposition 3.4.42. A consequence of this is that μsh_{V^σ} remains constant along a homotopy of contact embeddings. But, by Gromov's h-principle, Theorem 3.6.8, any two contact embeddings can be made homotopic to each other upon stabilizing, which we may do so that it does not change the sheaf on V . In summary, $\mu\text{sh}_{V, -\sigma}$ does not depend on f , but only, up to stabilization, on the homotopy class of σ . $\square$

Our next goal is to rephrase the above Lemma 3.6.12 in the language of stable homotopy theory and this will allow us to state a generalization of the construction. For this reason, we take a detour to discuss infinite loop spaces.

Definition 3.6.13. Equip the category of anima with the Cartesian symmetric monoidal structure $(\text{Ani}, \times)$. A monoid M in Ani is an algebra object $M \in \text{Alg}(\text{Ani}, \times)$, and we say that M is grouplike or a group object if the maps

$$\pi_1 \times m, \pi_2 \times m : M \times M \rightarrow M \times M$$

are equivalences where $m : M \times M \rightarrow M$ is the multiplication. Equivalently, M is a group object if the monoid in sets $\pi_0(M)$ is a group. We denote the category of group objects in Ani by $\text{Mon}^{\text{gp}}(\text{Ani})$ or, more precisely, $\text{Mon}_{\mathbb{E}_1}^{\text{gp}}(\text{Ani})$.

A natural source of group objects in anima is through taking the loop space $\Omega_*(X)$ for a pointed anima $X \in \text{Ani}_*$. In fact, taking iterated loop spaces increases the commutativity, as seen in classical algebraic topology by the fact that π_2 and above are abelian; in this higher categorical setting, the commutativity is controlled by the little disks operad $\mathbb{E}_k$. In this hierarchy, $\mathbb{E}_1$ controls associative algebras.

Theorem 3.6.14 ([45, Theorem 5.2.6.10]). *For $k > 0$, the functor of taking k -fold loop spaces $\Omega^k : \text{Ani}_* \rightarrow \text{Mon}_{\mathbb{E}_k}(\text{Ani})$ restricts to a fully faithful map on $\text{Ani}_*^{\geq k}$, and the image is given by the group-like $\mathbb{E}_k$ -monoids, $\text{Mon}_{\mathbb{E}_k}^{\text{gp}}(\text{Ani})$. Here, $\text{Ani}_*^{\geq k} \subseteq \text{Ani}_*$ consists of objects X such that $\pi_l(X) = *$ for $l < k$.³³*

We note that taking the loop space iteratively was discussed before, when we constructed the category of spectra Sp in Construction 2.4.11. The difference here is that we restrict ourselves, and this in fact recovers the subcategory of connective spectra Sp^{cn} mentioned in Remark 2.10.4.

³²More on the choice of notation $-\sigma$ right after this lemma.

³³We note the shift in notation from the usual n -connectedness. For example, 1-connective means path-connected, and 2-connective means simply connected, etc.

Corollary 3.6.15 ([45, Remark 5.2.6.26]). *The category $\mathrm{Sp}^{\mathrm{cn}}$ is given by the limit*

$$\mathrm{Sp}^{\mathrm{cn}} = \lim \left(\cdots \xrightarrow{\Omega} \mathrm{Ani}_*^{\geq 2} \xrightarrow{\Omega} \mathrm{Ani}_*^{\geq 1} \xrightarrow{\Omega} \mathrm{Ani}_* \right).$$

As a result, the equivalence in Theorem 3.6.14 assembles to an equivalence

$$\mathrm{Sp}^{\mathrm{cn}} = \mathrm{Mon}_{\mathbb{E}_\infty}^{\mathrm{gp}}(\mathcal{S}).$$

Consider the orthogonal group $O(n)$ and the unitary group $U(n)$ for some $n \geq 1$. As groups in topological spaces, they induce group objects in Ani so, by Theorem 3.6.14, there are classifying spaces $BO(n)$ and $BU(n)$. It can be shown that they classify real and complex vector bundles of dimension n . From this viewpoint, the inclusion map $O(n) \hookrightarrow U(n)$ induces, say for any manifold X , the map

$$(20) \quad \begin{aligned} \mathrm{Map}(X, BO(n)) &\rightarrow \mathrm{Map}(X, BU(n)) \\ E &\mapsto E \otimes_{\mathbb{R}} \mathbb{C} \end{aligned}$$

sending a real vector bundle $E \rightarrow X$ to its complexification.

Definition 3.6.16. For a complex vector bundle $\xi \in \mathrm{Map}(X, BU(n))$, a real vector bundle $\rho \in \mathrm{Map}(X, BO(n))$ is said to be a polarization of ξ if its image under (20) is equivalent to ξ .

Now, since there is an inclusion $U(n) \hookrightarrow \mathrm{Sp}(2n, \mathbb{R})$ into the real symplectic group, the automorphism group of linear transformations that preserve the standard symplectic vector space $(\mathbb{R}^{2n}, \omega_{st})$, any symplectic vector bundle can be endowed with a complex structure. Under this identification, a polarization ρ of a symplectic vector bundle ξ is equivalent to a Lagrangian subbundle of ξ . Since we assume the contact manifold V to be co-oriented, by Remark 3.6.2, it comes with a map

$$V \xrightarrow{\mathcal{H}} BU(n),$$

where $\dim V = 2n + 1$, classifying its contact distribution. To take into account the effect of stabilization, we consider the inclusions $O(n) \hookrightarrow O(n + 1)$ and $U(n) \hookrightarrow U(n + 1)$ at the upper left corner. This inclusion induces, for example, the map

$$(21) \quad \begin{aligned} \mathrm{Map}(X, BO(n)) &\rightarrow \mathrm{Map}(X, BO(n + 1)) \\ E &\mapsto E \oplus \mathbb{R}_X \end{aligned}$$

obtained by direct sum with the trivial line bundle. We then take

$$O := \operatorname{colim}_{n \rightarrow \infty} O(n) \text{ and } U := \operatorname{colim}_{n \rightarrow \infty} U(n),$$

and they classify real and complex vector bundles up to stabilization.

The virtue of stabilizing is that O and U are infinite-loop spaces, which is a consequence of Bott periodicity $\Omega^8 O = O$ and $\Omega^2 U = U$, and by Corollary 3.6.15, BO and BU are group objects themselves. Concretely, the multiplication is induced by passing

$$\begin{aligned} \mathrm{Map}(X, BO(n)) \times \mathrm{Map}(X, BO(m)) &\rightarrow \mathrm{Map}(X, BO(n + m)) \\ (E_1, E_2) &\mapsto E_1 \oplus E_2 \end{aligned}$$

to the colimit. The inverse is given by the “normal” vector bundle. That is, for any vector bundle $E \rightarrow X$, for some large enough $N \gg 1$, E can be embedded in a trivial bundle

$\mathbb{R}_X^N \rightarrow X$ over X , and the inverse of E is given by the quotient $\mathbb{R}_X^N/E$. This means that, in $\text{Map}(V, BU)$, the symplectic normal bundle $N(f)$ we consider in Lemma 3.6.12 is the inverse of the contact distribution $\mathcal{H}$. Since the morphism $BO \rightarrow BU$ is a group homomorphism, the Lagrangian subbundle σ that determines the thickening is the same as the inverse of the polarization $\rho = -\sigma$ of $\mathcal{H}$. Thus, we can rephrase Lemma 3.6.12 as follows:

Lemma 3.6.17. *Let V be a co-oriented contact manifold. A choice of (stable) polarization ρ of its contact structure gives a well-defined sheaf of categories, with the notation from Lemma 3.6.12, $\mu\text{sh}_{V;\rho} := \mu\text{sh}_{V;-\sigma}$ where $\sigma := -\rho$ in $\text{Map}(V, BO)$.*

Example 3.6.18. Let M be a C^∞ -manifold of dimension n . The fiber polarization ϕ_M is given by $\phi_M([x, \xi]) = T_x^*M$ for $[x, \xi] \in S^*M$.

In fact, the fiber polarization recovers the notion of microsheaves from the previous Section 3.5.

Proposition 3.6.19 ([11, Proposition 4.20]). *There is a canonical equivalence of sheaves of categories $\mu\text{sh}_{S^*M, \phi_M} = \mu\text{sh}_{S^*M}$ where the left-hand side is the definition given in Lemma 3.6.17 and the right-hand side is given in Definition 3.5.5.*

But, as a polarization is a homotopy invariant, on any contractible open subset $\mathcal{U}$, any two polarizations can be identified non-canonically. Thus,

Corollary 3.6.20. *Let V be a co-oriented contact manifold and ρ a polarization of its contact structure. Then, locally on a contractible Darboux chart $j : \mathcal{U} \hookrightarrow S^*M$, we have $\mu\text{sh}_{V;\rho}|_{\mathcal{U}} \cong j^*(\mu\text{sh}_{S^*M}|_{j(\mathcal{U})})$.*

The last discussion regarding the abstract theory is to upgrade the construction from polarizations to Maslov data: By Corollary 3.6.15, we know that the cofiber U/O is an infinite loop space. Alternatively, this follows again from Bott periodicity $\Omega(U/O) = \mathbb{Z} \times BO$ and $\Omega(\mathbb{Z} \times BO) = O$; in fact, $U/O = \text{colim}_{n \rightarrow \infty} U(n)/O(n)$ and $B(U/O)$ classifies stable Lagrangian Grassmannian bundles. Since taking $B(-)$ in $\text{Mon}_{\mathbb{E}_\infty}^{\text{gp}}(\mathbb{S})$ corresponds to taking $\Sigma(-)$ in Sp^{cn} , the fiber sequence $O \rightarrow U \rightarrow U/O$ induces a fiber sequence $BO \rightarrow BU \rightarrow B(U/O)$. Thus, a polarization of a (stable) complex vector bundle $\xi : X \rightarrow BU$ is nothing but a null homotopy of the composition

$$X \xrightarrow{\xi} BU \rightarrow B(U/O).$$

The reason to switch to this viewpoint is that we can further compose it with the twice delooping of the J-homomorphism $B(U/O) \rightarrow B^2 \text{Pic}(\mathbb{S})$. Here, for a symmetric monoidal stable category $(\mathcal{C}, \otimes)$, the Picard groupoid $\text{Pic}(\mathcal{C}) \subseteq \mathcal{C}$ is given by the non-full subcategory consisting of $\otimes$ -invertible objects and invertible morphisms. Since the coefficient category $\mathcal{V}$ we use is by definition a commutative algebra object $\text{Sp} \rightarrow \mathcal{V}$, there is an induced map $\text{Pic}(\mathbb{S}) \rightarrow \text{Pic}(\mathcal{V})$ by tensoring with $1_{\mathcal{V}}$.

Definition 3.6.21. Let V be a co-oriented contact manifold. A $(\mathcal{V})$ -Maslov datum τ on V is a null homotopy of the composition

$$V \xrightarrow{\mathcal{H}} BU \rightarrow B(U/O) \rightarrow B^2 \text{Pic}(\mathbb{S}) \rightarrow B^2 \text{Pic}(\mathcal{V}).$$

We note that a polarization canonically induces a Maslov datum by post-composition.

One of the main theorems in [50] is that the construction in Lemma 3.6.17 descends to Maslov data:

Theorem 3.6.22 ([50]). *Let V be a co-oriented contact manifold. A choice of Maslov datum τ gives a well-defined sheaf of categories $\mu\text{sh}_{V;\tau}$ locally given by μsh_{S^*M} on small Darboux charts.*

We end these notes by mentioning two appearances of microsheaves in related fields of mathematics. First, one can extend the definition of microsheaves to exact symplectic manifolds: Let (X, λ) be an exact symplectic manifold. Its contactization is the co-oriented contact manifold $(X \times \mathbb{R}, \lambda + ds)$. Since $\mathbb{R}$ is homotopically trivial, any polarization or Maslov datum extends uniquely from X to $X \times \mathbb{R}$.

Definition 3.6.23. The sheaf of microsheaves on X associated to a Maslov datum τ is the sheaf given by $\mu\text{sh}_{X;\tau} := \mu\text{sh}_{X \times \mathbb{R}, X \times \{0\};\tau} |_{X \times \{0\}}$.

An exact symplectic manifold (X, λ) is called Liouville if it has a cylindrical end, i.e. a controlled geometry at infinity modeled on a symplectization. In particular, it admits a notion of contact boundary $\partial_\infty X$ at infinity. It is called Weinstein if, in addition, it is equipped with an exhausting Morse function compatible with the Liouville vector field, providing a Morse-theoretic control of its dynamics. A consequence of being Weinstein is that the skeleton $\mathbf{c}_X$ defined to be the set of points that do not escape to infinity through the Liouville flow will be isotropic. Thus, if one chooses a (singular) Legendrian subset $\Lambda \subseteq \partial_\infty X$, the relative skeleton $\mathbf{c}_{X,\Lambda}$, given by $\mathbf{c}_X$ union the points flowing to Λ , will be isotropic. We remark that for a pair (X, Λ) , one can define the wrapped Fukaya category $\mathcal{W}(X, \Lambda)$ using Floer theory.

Theorem 3.6.24 ([26]). *Fix a polarization on a Weinstein manifold X and pick an isotropic set $\Lambda \subseteq \partial_\infty X$. Then, there is an equivalence*

$$\text{Ind}(\mathcal{W}(X, \Lambda)) = \mu\text{sh}_{\mathbf{c}_{X,\Lambda}}(\mathbf{c}_{X,\Lambda})$$

between the Ind-completion of the wrapped Fukaya category and the global sections of $\mu\text{sh}_{\mathbf{c}_{X,\Lambda}}$.

34

One can consider the same definition as in Definition 3.6.1 in the complex setting. In this case, the symplectization $\tilde{V}$ will be a $\mathbb{C}^\times$ -bundle over the original complex contact manifold V . As V is complex, it cannot have odd real dimension. However, the S^1 -bundle $\tilde{V}/\mathbb{R}^+ \rightarrow V$ will be real contact. In fact, one can show that there is a canonical Maslov datum τ_{can} on $\tilde{V}/\mathbb{R}^+$ coming from the contact structure. This thus induces a canonical sheaf $\mu\text{sh}_{\tilde{V}}$ by pulling back $\mu\text{sh}_{\tilde{V}/\mathbb{R}^+;\tau_{can}}$ to $\tilde{V}$ and one can ask whether the support of a section of $\mu\text{sh}_{\tilde{V}}$ is complex. Set $\mu\text{sh}_{V;\mathbb{C}-c} := \pi_* \mu\text{sh}_{\tilde{V};\mathbb{C}-c}$.

On the other hand, D-modules also admit a notion of microlocalization as E-modules. Slightly more precisely, let X be a complex manifold and $\mathcal{D}_X$ the sheaf of differential operators. There is a sheaf of microdifferential operators $\mathcal{E}_X$ on $\mathbb{P}^*X$ containing the pullback of $\mathcal{D}_X$ such that the extension is flat. Like the case of μsh , one can glue $\mathcal{E}_X$ canonically to an algebroid $\mathcal{E}_V$ on any complex contact manifold V .

Theorem 3.6.25 ([11, 12]).

- (1) *There is a canonical t -structure on $\mu\text{sh}_{V;\mathbb{C}-c}$, locally given by a microlocalization of the perverse t -structure.*

³⁴There is an ‘op’ in the cited paper and the difference is that we made a sign choice that is more natural for sheaf theory, and it happens to be the opposite of the convention there.

(2) *The Riemann-Hilbert correspondence extends to an equivalence*

$$(\mu\mathrm{sh}_{V;\mathbb{C}-c}^b)^\heartsuit = \mathcal{E}_V\text{-}\mathcal{M}\mathrm{od}_{rh}$$

between perverse microsheaves with perfect microstalks and regular holonomic $\mathcal{E}_V$ -modules.

Remark 3.6.26. As explained in [12], by taking the contactization, one can use the above microlocal Riemann-Hilbert correspondence to study W -modules with perverse microsheaves. W -modules are used to construct the notion of Category $\mathcal{O}$ for symplectic resolutions.

REFERENCES

- [1] Emmanuel Andronikof. A microlocal version of the Riemann-Hilbert correspondence. *Topol. Methods Nonlinear Anal.*, 4(2):417–425, 1994.
- [2] Tomohiro Asano and Yuichi Ike. Persistence-like distance on Tamarkin’s category and symplectic displacement energy. *J. Symplectic Geom.*, 18(3):613–649, 2020.
- [3] Tomohiro Asano and Yuichi Ike. The rectifiable rectangular peg problem. *arXiv preprint arXiv:2412.21057*, 2024.
- [4] Tomohiro Asano, Yuichi Ike, Christopher Kuo, and Wenyuan Li. C^0 -rigidity of legendrians and coisotropics via sheaf quantization. *arXiv:2510.01746*, 2025.
- [5] Qingyuan Bai and Yuxuan Hu. Toric mirror symmetry for homotopy theorists. *arXiv preprint arXiv:2501.06649*, 2025.
- [6] Alexander Beilinson and Joseph Bernstein. Localisation de $\mathfrak{g}$ -modules. *C. R. Acad. Sci., Paris, Sér. I*, 292:15–18, 1981.
- [7] Andrew J Blumberg, David Gepner, and Gonalo Tabuada. A universal characterization of higher algebraic k-theory. *Geometry & Topology*, 17(2):733–838, 2013.
- [8] Tom Braden, Anthony Licata, Nicholas Proudfoot, and Ben Webster. Quantizations of conical symplectic resolutions II: category $\mathcal{O}$ and symplectic duality. *arXiv preprint arXiv:1407.0964*, 2014.
- [9] Tom Braden, Nicholas Proudfoot, and Ben Webster. Quantizations of conical symplectic resolutions I: local and global structure. *Astérisque*, (384):1–73, 2016.
- [10] Jean-Luc Brylinski and Masaki Kashiwara. Kazhdan-lusztig conjecture and holonomic systems. *Inventiones mathematicae*, 64(3):387–410, 1981.
- [11] Laurent Côté, Christopher Kuo, David Nadler, and Vivek Shende. Perverse microsheaves. *arXiv:2209.12998*, 2022.
- [12] Laurent Côté, Christopher Kuo, David Nadler, and Vivek Shende. The microlocal Riemann-Hilbert correspondence for complex contact manifolds. *arXiv:2406.16222*, 2024.
- [13] Roger Casals and Wenyuan Li. Positive microlocal holonomies are globally regular. *arXiv:2409.07435*, 2024.
- [14] Kęstutis Česnavičius and Peter Scholze. Purity for flat cohomology. *Annals of Mathematics*, 199(1):51–180, 2024.
- [15] Sheng-Fu Chiu. Nonsqueezing property of contact balls. *Duke Mathematical Journal*, 166(4):605–655, 2017.
- [16] Bastiaan Cnossen, Tobias Lenz, and Sil Linskens. Universality of span 2-categories and the construction of 6-functor formalisms. *arXiv:2505.19192*, 2025.
- [17] Lee Cohn. Differential graded categories are k-linear stable infinity categories. *arXiv:1308.2587*, 2013.
- [18] Alexander I Efimov. K-theory and localizing invariants of large categories. *arXiv preprint arXiv:2405.12169*, 2024.
- [19] Yakov Eliashberg and Nikolai M Mishachev. *Introduction to the h-principle*. Number 48. American Mathematical Soc., 2002.
- [20] Bohan Fang, Chiu-Chu Melissa Liu, David Treumann, and Eric Zaslow. A categorification of Morelli’s theorem. *Invent. Math.*, 186(1):79–114, 2011.
- [21] Institute for Advanced Study and Univalent Foundations Program. *Homotopy Type Theory: Univalent Foundations of Mathematics*. Univalent Foundations Program, 2013.

- [22] Dennis Gaitsgory and Nick Rozenblyum. *A study in derived algebraic geometry. Vol. I. Correspondences and duality*, volume 221 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2017.
- [23] Benjamin Gammage and Vivek Shende. Mirror symmetry for very affine hypersurfaces. *arXiv:1707.02959v3*, 2021.
- [24] Sheel Ganatra, John Pardon, and Vivek Shende. Sectorial descent for wrapped Fukaya categories. *arXiv:1809.03427v2*, 2019.
- [25] Sheel Ganatra, John Pardon, and Vivek Shende. Covariantly functorial wrapped Floer theory on Liouville sectors. *Publ. Math. Inst. Hautes Études Sci.*, 131:73–200, 2020.
- [26] Sheel Ganatra, John Pardon, and Vivek Shende. Microlocal Morse theory of wrapped Fukaya categories. *arXiv:1809.08807v2*, 2020.
- [27] Stéphane Guillermou. Quantization of conic Lagrangian submanifolds of cotangent bundles. *arXiv:1212.5818*, 2012.
- [28] Stéphane Guillermou, Masaki Kashiwara, and Pierre Schapira. Sheaf quantization of Hamiltonian isotopies and applications to nondisplaceability problems. *Duke Math. J.*, 161(2):201–245, 2012.
- [29] Stéphane Guillermou and Claude Viterbo. The singular support of sheaves is γ -coisotropic. *Geometric and Functional Analysis*, 34(4):1052–1113, 2024.
- [30] Claudius Heyer and Lucas Mann. 6-functor formalisms and smooth representations. *arXiv:2410.13038*, 2024.
- [31] Ryoshi Hotta, Kiyoshi Takeuchi, and Toshiyuki Tanisaki. *D-modules, perverse sheaves, and representation theory*, volume 236 of *Progress in Mathematics*. Birkhäuser Boston, Inc., Boston, MA, 2008. Translated from the 1995 Japanese edition by Takeuchi.
- [32] Christian Houzel. A short history: Les débuts de la théorie des faisceaux. [36], pages 7–21. With a chapter in French by Christian Houzel, Corrected reprint of the 1990 original.
- [33] Marc Hoyois, Sarah Scherotzke, and Nicolò Sibilla. Higher traces, noncommutative motives, and the categorified Chern character. *arXiv:1511.03589v3*, 2017.
- [34] Masaki Kashiwara. The Riemann-Hilbert problem for holonomic systems. *Publications of the Research Institute for Mathematical Sciences*, 20(2):319–365, 1984.
- [35] Masaki Kashiwara and Takahiro Kawai. On holonomic systems of microdifferential equations, iii. systems with regular singularities. *Publications of the Research Institute for Mathematical Sciences*, 17(3):813–979, 1981.
- [36] 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.
- [37] Masaki Kashiwara and Pierre Schapira. Persistent homology and microlocal sheaf theory. *Journal of Applied and Computational Topology*, 2(1):83–113, 2018.
- [38] Maxim Kontsevich. Symplectic geometry of homological algebra. *Preprint. Available at https://www.ihes.fr/~maxim/TEXTS/Symplectic_AT2009.pdf*, 2009.
- [39] Christopher Kuo and Wenyuan Li. Duality, Künneth formulae, and integral transforms in microlocal geometry. *arXiv:2405.15211*, 2024.
- [40] Christopher Kuo and Wenyuan Li. Relative Calabi-Yau structure on microlocalization. *arXiv:2408.04085*, 2024.
- [41] Christopher Kuo and Wenyuan Li. Spherical adjunction and Serre functor from microlocalization. *arXiv:22210.06643*, 2024.
- [42] Christopher Kuo, Vivek Shende, and Bingyu Zhang. On the Hochschild cohomology of Tamarkin categories. *arXiv:2312.11447*, 2024.
- [43] Tatsuki Kuwagaki. The nonequivariant coherent-constructible correspondence for toric stacks. *arXiv:1610.03214*, 2017.
- [44] Jacob Lurie. *Higher topos theory*, volume 170 of *Annals of Mathematics Studies*. Princeton University Press, Princeton, NJ, 2009.
- [45] Jacob Lurie. Higher algebra. *Preprint, available at <https://www.math.ias.edu/~lurie/>*, 2017.
- [46] Jacob Lurie. Spectral algebraic geometry. *Preprint, available at <https://www.math.ias.edu/~lurie/>*, 2018.
- [47] Jacob Lurie. Kerodon. <https://kerodon.net>, 2026. Accessed: 2026-03-18.

- [48] Lucas Mann. A p-adic 6-functor formalism in rigid-analytic geometry. *arXiv preprint arXiv:2206.02022*, 2022.
- [49] David Nadler. Microlocal branes are constructible sheaves. *arXiv:math/0612399v4*, 2009.
- [50] David Nadler and Vivek Shende. Sheaf quantization in Weinstein symplectic manifolds. *arXiv:2007.10154v2*, 2021.
- [51] David Nadler and Eric Zaslow. Constructible sheaves and the Fukaya category. *J. Amer. Math. Soc.*, 22(1):233–286, 2009.
- [52] Andrei Okounkov. Lie algebras of the XXI century. Coloquio recording on YouTube, available at <https://www.youtube.com/watch?v=3ezeRLjQgI0>, August 2018.
- [53] Shaul Ragimov and Tomer M Schlank. The ∞ -categorical reflection theorem and applications. *Journal of Topology*, 19(1):e70060, 2026.
- [54] Emily Riehl and Dominic Verity. The 2-category theory of quasi-categories. *Advances in Mathematics*, 280:549–642, 2015.
- [55] Nick Rozenblyum. Filtered colimits of ∞ -categories. URL:<https://www.math.toronto.edu/nick/>, 2012.
- [56] Peter Scholze. Geometry and higher category theory. *Lecture Notes*, 2025.
- [57] Peter Scholze. Six-functor formalisms. *arXiv preprint arXiv:2510.26269*, 2025.
- [58] Peter Scholze. Wild betti sheaves. *arXiv preprint arXiv:2505.24599*, 2025.
- [59] Vivek Shende. Microlocal category for Weinstein manifolds via h-principle. *arXiv:1707.07663*, 2017.
- [60] Vivek Shende and Alex Takeda. Calabi-Yau structures on topological Fukaya categories. *arXiv:1605.02721v3*, 2020.
- [61] Germán Stefanich. Higher sheaf theory i: Correspondences. *arXiv:2011.03027*, 2020.
- [62] Dmitry Tamarkin. Microlocal condition for non-displaceability. In *Algebraic and analytic microlocal analysis. AAMA, Evanston, Illinois, USA, May 14–26, 2012 and May 20–24, 2013. Contributions of the workshops*, pages 99–223. Cham: Springer, 2018.
- [63] Bertrand Toën. Lectures on dg-categories. In *Topics in algebraic and topological K-theory*, volume 2008 of *Lecture Notes in Math.*, pages 243–302. Springer, Berlin, 2011.
- [64] Dmirty Vaintrob. The coherent-constructible correspondence and incomplete topologies. *Preprint available at the author’s webpage*, 2016.
- [65] Claude Viterbo. Sheaf quantization of Lagrangians and Floer cohomology. *arXiv:1901.09440*, 2019.
- [66] Marco Volpe. The six operations in topology. Preprint, arXiv:2110.10212 [math.AT] (2021), 2021.
- [67] Ingo Waschbies. The stack of microlocal perverse sheaves. *Bulletin de la société mathématique de France*, 132(3):397–462, 2004.
- [68] Bingyu Zhang. Capacities from the chiu-tamarkin complex. *arXiv:2103.05143v3*, 2022.
- [69] Peng Zhou. Twisted polytope sheaves and coherent-constructible correspondence for toric varieties. *arXiv:1701.00689*, 2017.