

A TOPOLOGICAL MODEL FOR THE FUKAYA CATEGORIES OF PLUMBINGS

MOHAMMED ABOUZAID

Abstract

We prove that the algebra of singular cochains on a smooth manifold, equipped with the cup product, is equivalent to the A_∞ structure on the Lagrangian Floer cochain group associated to the zero section in the cotangent bundle. More generally, given embeddings with isomorphic normal bundles of a closed manifold B into manifolds Q_1 and Q_2 , we construct a differential graded category from the singular cochains of these spaces, and prove that it is equivalent to the A_∞ category obtained by considering exact Lagrangian embeddings of Q_1 and Q_2 which intersect cleanly along B .

1. Introduction

In [8], Fukaya and Oh proved that counts of holomorphic discs in a cotangent bundle with boundary conditions along exact Lagrangian sections agree, in a certain degeneration, with a count of gradient trees. This has been widely expected to lead to an equivalence between the A_∞ structure defined on Floer cochains of a Lagrangian Q and the chain level cup product on the classical (Čech, simplicial, singular) models for its ordinary cohomology. Over \mathbb{R} , Kontsevich and Soibelman in [11] described an argument using the result of Fukaya and Oh and Homological perturbation theory on de Rham cohomology, which proves this result. A different proof, also over \mathbb{R} , is given by Fukaya, Oh, Ohta, and Ono's in their book on Lagrangian Floer cohomology [9, Section 33]. One corollary of this paper is that such an equivalence holds over the integers:

Theorem 1.1. *There is an A_∞ equivalence*

$$(1.1) \quad C^*(Q) \rightarrow CF^*(Q, Q).$$

REMARK 1.2. In order to define $CF^*(Q, Q)$ as an honest A_∞ algebra (rather than work with partially defined algebraic structures as in [11]) we follow the approach used by Seidel in [17]. The main idea is to make choices of perturbations to ensure the genericity of all moduli

Received 4/22/2009.

spaces that enter into the definition of the operations. Seidel proves that the resulting structure is independent, up to A_∞ equivalence, of these choices.

As in [11] we pass through a Morse model in order to prove this equivalence. First, we imitate the construction of the Fukaya category to obtain an A_∞ structure on Morse cochains by counting perturbed gradient trees. Writing $CM^*(Q)$ for this A_∞ algebra, we prove the existence of A_∞ functors

$$(1.2) \quad C^*(\mathcal{Q}) \rightarrow CM^*(Q) \leftarrow CF^*(Q, Q)$$

inducing isomorphisms on cohomology. Here, $C^*(\mathcal{Q})$ are the cochains of a simplicial triangulation of Q . Theorem 1.1 follows from the fact that A_∞ maps inducing an isomorphism on cohomology admit (quasi)-inverses whenever the underlying groups are free abelian.

We do not rely on the degeneration technique of [8] in order to prove the equivalence of the A_∞ structures coming from Floer and Morse theory. Rather, we consider moduli spaces that are built from gradient trees and holomorphic discs with one Lagrangian boundary condition along an arbitrary leaf of the foliation of the cotangent bundle by fibers. The corresponding abstract moduli spaces give a new realization of the multiplihedra controlling A_∞ functors foreseen by Stasheff in [19].

Theorem 1.1 is the special case of a more general result: within an exact symplectic manifold W with vanishing first chern class (and a choice of a complex volume form), we define a brane to be an exact Lagrangian Q which is relatively spin and such that the restriction of the complex volume form to some Weinstein neighborhood of Q is isotopic to the complexification of a (real) volume form on Q (more precisely, we fix this isotopy). These extra choices of data are necessary in order for the Fukaya category to be defined over \mathbb{Z} and admit natural gradings.

Given a pair of exact Lagrangian branes Q_1 and Q_2 intersecting cleanly along a submanifold B , the Lagrangian condition implies that the normal bundles of $Q_1 \cap Q_2$ in the two sheets are isomorphic. In Section 2.1, we construct a differential graded category

$$(1.3) \quad \text{Simp}(Q_1, Q_2)$$

with morphism spaces given by simplicial cochains on Q_i , NB , and $(NB, \partial NB)$ (here NB is the unit normal bundle of B in either Q_1 or Q_2). The main result of this paper relates this category to the full subcategory of the Fukaya category of W with objects Q_1 and Q_2 , and to a Morse theoretic model $\text{Morse}(Q_1, Q_2)$:

Theorem 1.3. *If ω vanishes on $\pi_2(M, Q_1 \cup Q_2)$, there are A_∞ equivalences*

$$(1.4) \quad \text{Simp}(Q_1, Q_2) \rightarrow \text{Morse}(Q_1, Q_2) \leftarrow \text{Fuk}(Q_1, Q_2).$$

REMARK 1.4. For clarity of exposition, we have chosen to write all arguments first whenever B is a point, with the case of clean intersection relegated to Appendix A. The proofs we give all extend straightforwardly to the clean intersection case, so the main point of the appendix is to explain the proper definitions and constructions in the general situation.

At the level of cohomology, Theorem 1.3 is not too difficult to prove from Poźniak’s thesis [15]. A standard application of Weinstein’s neighborhood theorem shows that a neighborhood of $Q_1 \cup Q_2$ in W is symplectomorphic to the result of gluing the cotangent bundles of Q_1 and Q_2 along the cotangent bundle of B . The diffeomorphism type of the resulting manifold is called the *plumbing* of the two cotangent bundles [14]. Up to an appropriate notion of deformation equivalence which does not affect the Fukaya category, there is a canonical symplectic form on this plumbing.

The experts will observe that the condition that ω vanish on $\pi_2(M, Q_1 \cup Q_2)$ is stronger than the mere exactness of Q_1 and Q_2 . One may think of it as a strict notion of exactness for the union of Q_1 and Q_2 as an immersed Lagrangian submanifold of M . For example, as we allow for B to be disconnected, our result gives a description of the category generated by a “cycle” of Lagrangians $\{Q_i\}_{i=1}^d$ such that Q_i intersects only Q_{i-1} and Q_{i+1} with the understanding that the labels are computed mod d . However, our description is only valid in the case we know in addition that a loop running around this cycle cannot bound any holomorphic discs (for energy reasons).

We recall that, in [16], Ruan extended the degeneration results of Fukaya and Oh to the case of an immersed Lagrangian $Q \in M$ with double points such that $\pi_2(M, Q)$ vanishes, and proved that counts of rigid holomorphic discs can be made to agree with counts of gradient trees. This indicates a straightforward generalization of Theorem 1.3: given two embeddings $B \rightarrow Q$ whose images do not intersect and with isomorphic normal bundles NB , it is easy to extend the construction of $\text{Simp}(Q_1, Q_2)$ to produce an associative product on the direct sum

$$(1.5) \quad C^*(Q) \oplus C^*(NB) \oplus C^*(NB, \partial NB)$$

which satisfies the Leibniz rule with respect to the obvious differential. This algebra is then equivalent, as an A_∞ algebra, to the Lagrangian Floer cochains of Q considered as an immersed Lagrangian in the “self-plumbing” of its cotangent bundle along B . Floer cohomology for immersed Lagrangians was described in great generality by Akaho and Joyce in [4], producing an A_∞ structure on a vector space which, up to choosing different model for a cochain-level theory, is equivalent to the direct sum (1.5). One may interpret our result as a computation, in the exact case, of the A_∞ structure assigned by Akaho and Joyce to

an immersed Lagrangian. Indeed, their A_∞ structure comes from the count of perturbed holomorphic discs; the perturbations are necessary because of the abundance of constant discs with marked points which do not form regular moduli spaces as soon as the number of marked points is greater than 3. As is often the case with such approaches, it seems impossible to give an explicit direct counts of the number of discs after perturbations.

Finally, we note that while every A_∞ category is abstractly equivalent to a differential graded category, there is no construction which produces a finite dimensional differential graded category when the morphism spaces are finite rank abelian groups. One consequence of Theorem 1.3 is the existence of such a finite dimensional differential graded model for Fukaya categories in the case of plumbings. We insist, however, that our construction requires working with cochains on Q_1 and Q_2 even if these manifolds happen to be formal (i.e., have the property that their cohomology are quasi-isomorphic to their cochain algebras). As explained in Chapter 20 of [17], it is already true that one can find 5 exact Lagrangians in a twice-punctured genus 2 curve, satisfying our exactness condition, such that the subcategory of the Fukaya category they generate is not formal (despite the fact that the circle is formal).

Acknowledgments. This research was conducted during the period the author served as a Clay Research Fellow. He would like to thank Paul Seidel for answering many questions about Chapter 20 of [17].

2. Three categories

For the remainder of this paper, we fix a pair of closed smooth manifolds Q_1 and Q_2 with base points b_i , and open sets U_i containing b_i . In the most basic situation, U_i is the interior of a closed (embedded) ball. We shall write $Q_{1,2}$ for the closure of U_1 and $Q_{2,1}$ for the closure of U_2 .

In addition, we equip Q_1 and Q_2 with simplicial triangulations \mathcal{Q}_1 and \mathcal{Q}_2 , and choose top-dimensional simplices Δ_1 and Δ_2 , respectively containing b_1 and b_2 . Recall that the condition that a triangulation be simplicial implies that we have an ordering of all vertices; in particular, Δ_i has a unique identification ι_i with the standard n -simplex Δ which preserves the ordering (see pages 102 and 107 of [10] for a discussion of different notions of triangulations).

2.1. A simplicial category from a triangulation of Q_i . We begin by constructing a differential graded category with two objects called Q_1 and Q_2 and morphisms the simplicial cochain complexes

$$(2.1) \quad \mathrm{Hom}_*^S(Q_i, Q_i) = C^*(\mathcal{Q}_i),$$

$$(2.2) \quad \mathrm{Hom}_*^S(Q_1, Q_2) = C^*(\Delta),$$

$$(2.3) \quad \mathrm{Hom}_*^S(Q_2, Q_1) = C^*(\Delta, \partial\Delta).$$

In particular, the rank of $\text{Hom}(Q_1, Q_2)$ is $2^n - 1$, while the rank of $\text{Hom}(Q_2, Q_1)$ is 1. The cochain complexes $C^*(Q_i)$ are differential graded algebras with respect to the cup product; we shall extend this product to a differential graded composition law by noting that the inclusions of Δ_i in Q_i induce ring homomorphisms

$$(2.4) \quad \iota_i^*: C^*(Q_i) \rightarrow C^*(\Delta).$$

We define compositions using the cup product on the left and on the right with the restricted classes

$$(2.5) \quad C^*(\Delta) \otimes C^*(Q_1) \xrightarrow{\text{id} \otimes \iota_1^*} C^*(\Delta) \otimes C^*(\Delta) \xrightarrow{\cup} C^*(\Delta),$$

$$(2.6) \quad C^*(Q_2) \otimes C^*(\Delta) \xrightarrow{\iota_2^* \otimes \text{id}} C^*(\Delta) \otimes C^*(\Delta) \xrightarrow{\cup} C^*(\Delta).$$

Since the maps ι_i^* are maps of differential graded algebras, and the Leibniz and associativity formulae hold for the cup product and differential on $C^*(\Delta)$, we readily conclude that the multiplication (2.5) satisfies the Leibniz rule, and moreover defines an associative composition. For example, we have the following commutative diagram

$$(2.7) \quad \begin{array}{ccc} C^*(Q_2) \otimes C^*(\Delta) \otimes C^*(Q_1) & \xrightarrow{\quad\quad\quad} & C^*(\Delta) \otimes C^*(\Delta) \otimes C^*(Q_1) \\ \downarrow & \searrow & \swarrow \\ & C^*(\Delta) \otimes C^*(\Delta) \otimes C^*(\Delta) & \\ \swarrow & \nearrow & \downarrow \\ C^*(Q_2) \otimes C^*(\Delta) \otimes C^*(\Delta) & \xrightarrow{\quad\quad\quad} & C^*(\Delta) \end{array}$$

In order to define this composition in the other direction, we use the fact that the cup product also makes $C^*(\Delta, \partial\Delta)$ into a bimodule over $C^*(\Delta)$. In particular, we define composition maps

$$(2.8) \quad C^*(Q_1) \otimes C^*(\Delta, \partial\Delta) \xrightarrow{\iota_1^* \otimes \text{id}} C^*(\Delta) \otimes C^*(\Delta, \partial\Delta) \xrightarrow{\cup} C^*(\Delta, \partial\Delta),$$

$$(2.9) \quad C^*(\Delta, \partial\Delta) \otimes C^*(Q_2) \xrightarrow{\text{id} \otimes \iota_2^*} C^*(\Delta, \partial\Delta) \otimes C^*(\Delta) \xrightarrow{\cup} C^*(\Delta, \partial\Delta).$$

Again, the proof that these maps satisfy the properties required of a differential graded category follows immediately from the fact that $C^*(\Delta, \partial\Delta)$ is in fact a differential graded bi-module over $C^*(\Delta)$ (this is trivial since the differential on this relative cochain group vanishes). Finally, the inclusions of Δ_i in Q_i induce maps

$$(2.10) \quad \iota_{i!}: C^*(\Delta, \partial\Delta) \rightarrow C^*(Q_i).$$

Since $C^*(\Delta, \partial\Delta)$ has rank 1, there is an obvious formula for these maps; the image is simply the n -dimensional cochain with value 1 on Δ_i and

which vanishes on every other chain. More conceptually, this map is induced by a composition

$$(2.11) \quad C^*(\Delta, \partial\Delta) \xrightarrow{\sim} C^*(Q_i, Q_i - \text{int}(\Delta_i)) \rightarrow C^*(Q_i),$$

where the first map is an excision isomorphism. Using this, we define the final two composition maps

$$(2.12) \quad C^*(\Delta, \partial\Delta) \otimes C^*(\Delta) \xrightarrow{\cup} C^*(\Delta, \partial\Delta) \xrightarrow{\iota_1!} C^*(Q_1),$$

$$(2.13) \quad C^*(\Delta) \otimes C^*(\Delta, \partial\Delta) \xrightarrow{\cup} C^*(\Delta, \partial\Delta) \xrightarrow{\iota_2!} C^*(Q_2).$$

There are eight associativity diagrams whose commutativity remains to be checked. Even though the construction is not completely symmetric, we only discuss the following four maps:

$$(2.14) \quad \begin{aligned} C^*(Q_2) \otimes C^*(\Delta) \otimes C^*(\Delta, \partial\Delta) &\longrightarrow C^*(Q_2), \\ C^*(\Delta) \otimes C^*(Q_1) \otimes C^*(\Delta, \partial\Delta) &\longrightarrow C^*(Q_2), \\ C^*(\Delta) \otimes C^*(\Delta, \partial\Delta) \otimes C^*(Q_1) &\longrightarrow C^*(Q_2), \\ C^*(\Delta, \partial\Delta) \otimes C^*(\Delta) \otimes C^*(\Delta, \partial\Delta) &\longrightarrow C^*(\Delta). \end{aligned}$$

Associativity of the first three of the above operations follows immediately from the property that $C^*(\Delta, \partial\Delta)$ is a differential graded bimodule over $C^*(\Delta)$. The last operation requires in addition the fact that the composition

$$(2.15) \quad C^*(\Delta, \partial\Delta) \xrightarrow{\iota_1!} C^*(Q_i) \xrightarrow{\iota_i^*} C^*(\Delta)$$

agrees with the natural inclusion of relative cochains into ordinary cochains (alternatively, one can notice that the two maps we are comparing in Equation (2.14) vanish altogether). We have provided all the necessarily ingredients to construct a differential graded category:

Definition 2.1. *The morphisms in Equation (2.1) and the composition laws given by Equations (2.5), (2.8), and (2.12) satisfy the axioms of a differential graded category. We denote by $\text{Simp}(Q_1, Q_2)$ the corresponding A_∞ category where the differential and the composition are twisted by a sign*

$$(2.16) \quad \mu_1^S(\check{\sigma}) = (-1)^{\deg(\check{\sigma})} d\check{\sigma},$$

$$(2.17) \quad \mu_2^S(\check{\sigma}, \check{\tau}) = (-1)^{\deg(\check{\tau})} \check{\sigma} \cup \check{\tau}.$$

We shall also need a geometric interpretation of the cup product as an intersection product on the chains of subdivisions dual to the chosen triangulations of Q_i . Let us therefore fix such dual subdivisions \check{Q}_i , whose top-dimensional cells are identified, by smooth charts, with polyhedra in \mathbb{R}^n . In particular, the $(n-1)$ -dimensional cells intersect cleanly.

REMARK 2.2. The sign change in (2.16) is particularly convenient from the point of view of the dual cell subdivision: if we think geometrically, every generator of $C_*(\mathcal{Q})$ corresponds to a simplex carrying an orientation, and the differential is the restriction of this orientation to the boundary. If we want to think of the differential d on $C^*(\mathcal{Q})$ as taking a dual cell to its boundary, we find that the resulting operation differs from the natural restriction of orientations exactly by a sign $(-1)^{1+\deg(\check{\sigma})}$. The additional sign of 1 comes from the fact that an outward pointing vector at the boundary of a cell corresponds to an inward-pointing vector from the dual subdivision's point of view.

We shall also require that the identification $\Delta_1 \cong \Delta_2$ map the cells of $\check{\mathcal{Q}}_1 \cap \Delta_1$ to the corresponding cells of $\check{\mathcal{Q}}_2 \cap \Delta_2$. In particular, we obtain a subdivision dual to the trivial triangulation of the standard simplex Δ . Given a cell σ of \mathcal{Q}_i , we write $\check{\sigma}$ for the dual cell of \mathcal{Q}_i , as well as for the dual generator of $C^*(\mathcal{Q}_i)$. A slight generalization of the main observation in Appendix E of [2] is the existence of vector fields X_i on \mathcal{Q}_i which generate flows ϕ_t^i such that the following Lemma holds

Lemma 2.3. *Whenever σ , τ , and ρ are cells of \mathcal{Q}_i such that*

$$(2.18) \quad \check{\sigma} \cup \check{\tau} = \check{\rho},$$

there is a smooth map

$$(2.19) \quad \check{\rho} \times [0, 1] \rightarrow \mathcal{Q}_i$$

such that the image of $\check{\rho} \times \{t\}$ agrees with $\check{\sigma} \cap \phi_t^i(\check{\tau})$ whenever t is sufficiently small. q.e.d.

Since $\check{\sigma} \cup \check{\tau}$ either vanishes or is of the form $\check{\rho}$ for some cell ρ , we conclude that the cup product $\check{\sigma} \cup \check{\tau}$ vanishes whenever $\check{\sigma} \cap \phi_t^i(\check{\tau})$ is empty for sufficiently small t , and we interpret the above lemma to state that the cup product is given by the intersection product of the dual cells, provided that they are appropriately perturbed.

The construction of the vector fields X_i is elementary; since \mathcal{Q}_i is a simplicial triangulation, its vertices are ordered, and hence so are the top-dimensional cells of dual subdivision. We may choose X_i to be any vector field satisfying the following property.

Definition 2.4. *A vector field X_i is said to be **compatible with the simplicial triangulation** if along each cell $\check{\sigma}$ the restriction of X_i lies in the tangent space of the top-dimensional cell which is minimal among those adjacent to σ .*

By abuse of language, we shall often call such vector fields **simplicial**. The existence of such a vector field follows from a local argument described in Appendix E of [2]. Since the above condition is respected by taking positive linear combinations, we may and shall require the restrictions of X_1 and X_2 to Δ_1 and Δ_2 to be intertwined by the maps




Figure 1

identifying these two simplices. Figure 1 shows the restriction of such a vector field to the 1-skeleton of the dual subdivision of a 2 dimensional simplex.

In terms of the dual subdivisions \check{Q}_i , the restriction of cochains corresponds to taking the inverse image of a cell, while the map $\iota_{i!}$ takes the vertex of the subdivision dual to Δ to its image under the inclusions into Q_i . Since composition maps in the category $\text{Simp}(\mathcal{Q}_1, \mathcal{Q}_2)$ are defined using these operations together with the cup product, Lemma 2.3 provides an intersection theoretic interpretation of this category.

2.2. A Morse category. Fix a pair of open sets $U_i \subset Q_i$ together with diffeomorphisms

$$(2.20) \quad \bar{U}_1 \cong Q_{1,2} \cong Q_{2,1} \cong \bar{U}_2,$$

where $Q_{1,2}$ and $Q_{2,1}$ are manifolds with boundary (there is no gain in notational simplicity in assuming anything more). To make the connection with the rest of the paper, we shall later restrict to the situation where U_1 and U_2 are both balls, corresponding to neighborhoods in Q_1 and Q_2 of the point b along which we shall perform the plumbing.


The identification between these manifolds shall be fixed throughout the following discussion, and we use it to define the space

$$Q = Q_1 \cup_{U_1 \cong U_2} Q_2$$

which is a non-Hausdorff smooth manifold whose locus of points which cannot be separated from each other is the union $\partial\bar{U}_1 \cup \partial\bar{U}_2$. By construction, this space admits canonical maps from the manifolds $Q_{1,2}$ and $Q_{2,1}$ as well. For the remainder of the paper, we sometimes also write $Q_{i,i}$ for Q_i as a convenient notational device.

We assume that we have chosen Riemannian metrics on Q_i whose restrictions to U_i are intertwined by these diffeomorphisms, and we fix a choice of Morse-Smale functions

$$(2.21) \quad f_{i,j}: Q_{i,j} \rightarrow \mathbb{R}$$


Figure 2

such that $f_{1,2}$ and $f_{2,1}$ are, respectively, inward and outward pointing on the boundary. In particular, we also choose functions $f_{1,1}$ and $f_{2,2}$ on Q_1 and Q_2 . Writing $CM^*(f)$ for Morse cochains, we shall define an A_∞ category with morphism space

$$(2.22) \quad Hom_*^M(Q_i, Q_j) = CM^*(f_{i,j}) = \bigoplus_{x \in \text{Crit}(f_{i,j})} |o_x|$$

where $|o_x|$ is the orientation line associated to $o_x \equiv \lambda(W^s(x))$, i.e. the quotient of the free abelian group generated by a choice of orientation of the stable manifold of x by the relation that the sum of generators associated to opposite orientations vanishes. This rank 1 group has degree equal to the dimension of the stable manifold of x . The differential counts negative gradient lines of $f_{i,j}$; the detailed description is given in Section 8.

To define higher compositions in this category, we consider moduli spaces of gradient trees in Q which are defined later in this section. Given any ribbon tree T with $d+1$ leaves, we shall assume that one of the leaves has been distinguished as outgoing, which determines the unique ordering on the remaining d leaves (called incoming) illustrated in Figure 2. We shall write $\mathcal{E}(T)$ for the set of edges of T (including the external edges). Our conventions are that all edges include their endpoints and, except for the outgoing edge, which is always infinite, are allowed to have length 0.

The edges adjacent to any vertex of a ribbon tree T are by definition cyclically ordered, and we define a labeling $(i(e), j(e))$ of the edges $e \in \mathcal{E}(T)$ by elements of the set $\{1, 2\}$ to be compatible with the cyclic ordering if $i(e') = j(e)$ whenever e' follows e with respect to the cyclic ordering at some vertex. The following result may be proved by induction or by embedding T in a disc labeling the complementary regions appropriately.

Lemma 2.5. *A sequence $\vec{I} = (i_0, \dots, i_d)$ determines a unique labeling of all edges of T which is compatible with the cyclic ordering, such that*

the incoming leaves are labeled by (i_k, i_{k+1}) . Moreover, the outgoing leaf is labeled by (i_0, i_d) . q.e.d.

In order to resolve potential problems with transversality in defining the Morse category, we introduce perturbation data, following Seidel's construction for Lagrangian Floer theory in [17]. Let us illustrate the case of the cup product on $CM^*(Q_1)$. Given three critical points x_0 , x_1 , and x_2 , the x_0 component of $\mu_2^M([x_2], [x_1])$ should count triple intersections between the ascending manifold of x_0 and the descending manifolds of x_1 and x_2 . This space is necessarily empty unless $x_1 = x_2$, in which case the intersection is not transverse.

This triple intersection admits a description as the moduli space of maps from the unique trivalent tree with three infinite external edges, which converge at the respective ends to the critical points x_0 , x_1 , and x_2 , and which on each edge solve the gradient flow equation for the function $f_{1,1}$. We shall perturb this moduli space by changing the differential equation obeyed by each edge. Concretely, as each edge is isometric to $[0, +\infty)$, our perturbation will vanish on $(1, +\infty)$ and will be essentially arbitrary on $[0, 1]$. Integrating the perturbed gradient flow along the interval $[0, 1]$, we obtain a triple ϕ_0 , ϕ_1 , and ϕ_2 of diffeomorphisms of Q_1 . Changing our perspective once again, we interpret the moduli space of perturbed gradient trees as the intersection of the images under ϕ_i of the appropriate ascending and descending manifolds. As the diffeomorphisms ϕ_i are essentially arbitrary (and independent of each other), we achieve transversality (and hence an honest count) upon perturbing the gradient flow. One must still prove that this product is the quadratic term of an A_∞ structure, which requires defining compatible perturbations of the gradient flow equation on higher dimensional moduli spaces of trees.

Let us return to the general discussion and fix a sequence of labels \vec{I} as above:

Definition 2.6. *A gradient flow perturbation datum on (T, g_T) is a choice, for each edge $e \in \mathcal{E}(T)$, of a family of vector fields*

$$(2.23) \quad X_e: e \rightarrow C^\infty(TQ_{i(e),j(e)})$$

which vanish away from a bounded subset of e and such that the restriction of the vector field

$$(2.24) \quad -\nabla f_{i(e),j(e)} + X_e$$

to the boundary of $Q_{i(e),j(e)}$ is (i) outward pointing if $(i(e), j(e)) = (1, 2)$ and (ii) inward pointing if $(i(e), j(e)) = (2, 1)$.

REMARK 2.7. Since the gradient vector fields of $f_{1,2}$ and $f_{2,1}$ are, respectively, inward and outward pointing, the vanishing perturbation datum satisfies the above properties.

We shall now specialize and consider the trees which control the A_∞ structure on the Morse category. Let us write \mathcal{T}_d for the moduli space of **Stasheff trees** (T, g_T) , i.e., metric ribbon trees with $d + 1$ infinite external edges, and $\mathcal{T}_{\vec{I}}$ for the moduli space of trees whose inputs are labeled by the successive elements of a sequence \vec{I} (this is a copy of $\mathcal{T}_{|\vec{I}|}$). By allowing singular trees, we obtain compactifications $\overline{\mathcal{T}}_d$ and $\overline{\mathcal{T}}_{\vec{I}}$ which are in fact polyhedra. The collection of spaces $\overline{\mathcal{T}}_d$ for varying d were shown by Stasheff to form an A_∞ operad, i.e., to control A_∞ algebras. In particular whenever $d = d_1 + d_2 - 1$, we may construct a singular Stasheff tree with d inputs by attaching the output of a singular tree with d_2 inputs to any of the inputs of an element of $\overline{\mathcal{T}}_{d_1}$. This construction yields a collection of d_1 maps

$$(2.25) \quad \overline{\mathcal{T}}_{d_1} \times \overline{\mathcal{T}}_{d_2} \rightarrow \overline{\mathcal{T}}_d$$

whose images for all pairs (d_1, d_2) cover the boundary of $\overline{\mathcal{T}}_d$. This familiar construction can also be done with labels: if $\vec{I}^1 = (i_0^1, \dots, i_{d_1}^1)$ and $\vec{I}^2 = (i_0^2, \dots, i_{d_2}^2)$ are sequences such that (i_k^1, i_{k+1}^1) is the label for the output of a tree in $\mathcal{T}_{\vec{I}^d}$, we obtain a natural map

$$(2.26) \quad \overline{\mathcal{T}}_{\vec{I}^1} \times \overline{\mathcal{T}}_{\vec{I}^2} \rightarrow \overline{\mathcal{T}}_{\vec{I}},$$

where $\vec{I} = (i_0^1, \dots, i_k^1 = i_0^2, i_1^2, \dots, i_{d_2}^2 = i_{k+1}^1, \dots, i_{d_1}^1)$ and the boundary of $\overline{\mathcal{T}}_{\vec{I}}$ is again covered by the images of such maps.

Definition 2.8. *A universal consistent perturbation datum for trees is a choice \mathbf{X}^T of a smooth family of perturbation data for elements of $\overline{\mathcal{T}}_{\vec{I}}$ for every sequence \vec{I} , which is compatible with the gluing maps (2.26) and is invariant under the automorphisms of each tree.*

REMARK 2.9. The meaning of a smoothness is clarified in Definition 7.1. A priori, each edge in a singular ribbon tree is equipped with at least two perturbation data: one comes from the singular tree, and the other from the (smooth) tree wherein the edge lies. Compatibility with the maps (2.26) is the requirement that these two perturbation data agree. The condition of invariance implies that the perturbation datum vanishes whenever there is only one input (because the \mathbb{R} translation symmetry forces a non-vanishing perturbation datum to have non-compact support, contradicting the boundedness requirement in Definition 2.6).

Every edge e of a Stasheff tree is isometric to a segment in \mathbb{R} . Writing t_e for the induced coordinate, we define a (perturbed) gradient flow line on e to be a map ψ_e from e to $Q_{i(e), j(e)}$ which is a solution to the differential equation

$$(2.27) \quad d\phi_e(\partial_{t_e}) = -\nabla f_{i(e), j(e)} + X_e.$$

Definition 2.10. A perturbed gradient tree ψ on (T, g_T) with labels \vec{I} is a continuous map

$$(2.28) \quad \psi: T \rightarrow Q$$

whose restriction to every edge $e \in \mathcal{E}(T)$ lifts to a map

$$(2.29) \quad \psi_e: e \rightarrow Q_{i(e), j(e)}$$

which is a gradient trajectory of $f_{i(e), j(e)}$.

REMARK 2.11. One usually imposes a balancing condition at vertices of T ; i.e., the sum of all vector fields associated to edges adjacent to v is required to vanish. This constraint is not part of our setup for gradient trees, and in fact our choices of perturbation data are such that we can choose the vector fields defined by the various edges which share a common vertex to have arbitrary and independent values, as long as the conditions on the boundary of $Q_{1,2}$ or $Q_{2,1}$, described in Definition 2.6, are satisfied.

If $\vec{x} \equiv (x_1, \dots, x_d)$ is a sequence of critical points of the functions $f_{i_k, i_{k+1}}$, and x_0 is a critical point of f_{i_0, i_d} , define

$$(2.30) \quad \mathcal{T}_d(x_0, \vec{x})$$

to be the set of isomorphism classes of perturbed gradient trees ψ such that the image of the k th incoming leaf is x_k , and the image of the unique outgoing leaf is x_0 . This moduli space admits a bordification $\overline{\mathcal{T}}_d(x_0, \vec{x})$ by adding all singular gradient trees.

We begin by showing that these moduli spaces are compact:

Lemma 2.12. *Given x_0 a critical point of f_{i_0, i_d} , and (x_1, \dots, x_d) critical points of the functions $f_{i_k, i_{k+1}}$, the moduli space $\overline{\mathcal{T}}(x_0, \vec{x})$ is compact.*

Sketch of proof. When studying Morse theory on a manifold with boundary, one must first show that gradient trees cannot escape to the boundary:

CLAIM. There exists a neighborhood of the boundary of $Q_{i,j}$ whose image in Q is disjoint from the image of every edge e labeled by $Q_{i,j}$ under a gradient tree

$$(2.31) \quad \psi: (T, g_T) \rightarrow Q.$$

Proof of claim. The assumption that the perturbation data are consistent implies that there is a uniform (i.e., independent of (T, g_T) in a fixed space \mathcal{T}_d) neighborhood of the boundary $Q_{1,2}$ (respectively, $Q_{2,1}$) in which the perturbed gradient flow (2.24) associated to every edge labeled by $(1, 2)$ (or $(2, 1)$) of a tree in \mathcal{T}_d is inward (respectively, outward) pointing in the sense that it increases the distance to the boundary. In particular, the gradient flows of $f_{1,2}$ and $f_{2,1}$ are themselves inward or

outward pointing in this neighborhood. By shrinking these neighborhoods, we may assume that they are identified by the fixed diffeomorphism $Q_{1,2} \cong Q_{2,1}$. We shall prove that the images of edges labeled by $Q_{1,2}$ or $Q_{2,1}$ cannot intersect this neighborhood.

The case of edges labeled by $Q_{2,1}$ is simplest. Note that if e is such an edge, then it must lie on a descending arc $(e_0, \dots, e_r = e)$ such that e_0 is an incoming leaf and each edge of the arc is labeled by $Q_{2,1}$. Assuming by contradiction that the image of e intersects the neighborhood of the boundary where the perturbed negative gradient flow points inward, we conclude that e_0 is contained in this fixed neighborhood. This implies that the image of this incoming leaf is a critical point which is also contained in this neighborhood, which contradicts the fact that $f_{2,1}$ has no critical points near the boundary.

If an edge is labeled by $Q_{1,2}$, there are two possibilities: either (i) there is a descending arc (e, e_1, \dots, e_r) with e_r the outgoing leaf such that all edges succeeding e are labeled by $Q_{1,2}$ or (ii) there exists a descending arc $(e = e_0, e_1, \dots, e_r)$ all of whose edge are labeled by $Q_{1,2}$ and an ascending arc $(e' = e'_0, \dots, e'_s)$ whose edges are labeled by $Q_{2,1}$ such that e_r and e' are adjacent. Using the same argument as in the previous case, we conclude that the critical point of $f_{1,2}$ (or $f_{2,1}$) whither e_r (or e'_s) limits must be contained in the previously fixed neighborhood of the boundary of $Q_{2,1}$ (or $Q_{1,2}$) contradicting our assumptions that the gradient flow may not vanish in these regions. q.e.d.

This implies that whenever the image of an edge under a family of gradient trees converges toward the singularities of Q , the edge must be labeled by $Q_{1,1}$ or $Q_{2,2}$, so all analysis can be done locally in either manifold. The standard proof of compactness for gradient trees for compact smooth Riemannian manifolds can then be used to prove the desired result (see [8]). q.e.d.

The following is the main result proved in Section 7.1:

Lemma 2.13. *For a generic choice of universal perturbation data, all spaces $\overline{\mathcal{T}}_d(x_0, \vec{x})$ are naturally compact topological manifolds with boundary of dimension*

$$(2.32) \quad d - 2 + \deg(x_0) - \sum_{1 \leq k \leq d} \deg(x_k).$$

REMARK 2.14. Technically, Q is a smooth non-Hausdorff manifold, with the set of points which cannot be separated from each other diffeomorphic to $\partial Q_{1,2}$. Any local construction for smooth manifolds is valid on Q as long as it performed away from this subset. In addition, if we are working near this subset, and we have additional data that distinguishes one of the sheets, then we may again use standard constructions from the study of smooth manifolds. Note that this is precisely what happens

if we are considering an edge of a gradient tree labeled by (i, i) whose image in Q happens to intersect the non-Hausdorff locus; by definition, we have a lift to the relevant (usual) manifold Q_i . The key component in the proof of Lemma 2.13 is the Claim appearing in Lemma 2.12 which asserts that there is a neighborhood of the non-Hausdorff locus of Q which (perturbed) gradient edges labeled by $Q_{i,j}$ with $i \neq j$ cannot intersect if they form part of a gradient Stasheff tree.

Given a sequence \vec{x} such that

$$(2.33) \quad d - 2 + \deg(x_0) - \sum_{1 \leq k \leq d} \deg(x_k) = 0,$$

we conclude that the moduli space $\mathcal{T}_d(x_0, \vec{x})$ consists of finitely many points. The description of this moduli space as a fiber product determines an isomorphism

$$(2.34) \quad \lambda(\mathcal{T}(x_0, \vec{x})) \otimes \lambda(Q^{d+1}) \cong \lambda(Q) \otimes \lambda(\mathcal{T}_d) \otimes \lambda(W^s(x_0)) \otimes \lambda(W^u(\vec{x}))$$

where $W^u(\vec{x})$ is the product of the descending manifolds of the critical points x_k . Our conventions are explained in Section 8. Whenever ψ is a rigid tree in $\mathcal{T}(x_0, \vec{x})$, the above isomorphism and Equation (8.4) give a natural map

$$(2.35) \quad \mathfrak{o}_{x_d} \otimes \cdots \otimes \mathfrak{o}_{x_1} \rightarrow \mathfrak{o}_{x_0}.$$

We define the d th higher product

$$(2.36) \quad \mu_d^M: CM^*(f_{i_{d-1}, i_d}) \otimes \cdots \otimes CM^*(f_{i_1, i_2}) \rightarrow CM^*(f_{i_1, i_d})$$

to be a sum over the induced maps μ_ψ on orientation bundles

$$(2.37) \quad [x_d] \otimes \cdots \otimes [x_1] \\ \mapsto \sum_{\substack{\psi \in \mathcal{T}(x_0, \vec{x}) \\ x_0}} (-1)^{(n+1)(\deg(x_0) + \dagger(\vec{x}))} \mu_\psi([x_d] \otimes \cdots \otimes [x_1])$$

where the sign is given by

$$(2.38) \quad \dagger(\vec{x}) = \sum_{k=1}^d k \deg(x_k).$$

The following result follows directly from Proposition 2.13 and the proof of the analogous result in Appendix C of [2].

Proposition 2.15. *The operations μ_d^M satisfy the axioms of an A_∞ category. q.e.d.*

In particular, we define the A_∞ category $\text{Morse}(Q_1, Q_2)$ to have objects and morphisms as in Equation (2.22), and (higher) compositions defined by the operations μ_d^M .

REMARK 2.16. The orientation convention used in [2] differs in a very minor way from the one introduced here. In particular, Equation (2.38) is a simplification of Equation (C.2) in the previous paper.

2.3. The Fukaya category of the skeleton of a plumbing. If \mathbb{C}^n is equipped with the standard symplectic form, complex structure, and with the coordinates $(x_1, \dots, x_n, y_1, \dots, y_n) = (\vec{x}, \vec{y})$, the Lagrangians

$$\begin{aligned} L_1 &= \mathbb{R}^n = \{\vec{y} = 0\}, \\ L_2 &= i\mathbb{R}^n = \{\vec{x} = 0\} \end{aligned}$$

intersect transversely at the origin. We define the plumbing

$$(2.39) \quad D^*L_1 \# D^*L_2$$

to be the open neighborhood of radius 1 of the union of L_1 and L_2 (in the standard euclidean metric).

We shall now use this local model to construct a Liouville manifold by gluing cotangent bundles: pick a Riemannian metric on Q_i whose restriction to a neighborhood of b_i is isometric to the ball of radius 4 in \mathbb{R}^n with b_i mapping to the origin and U_i to the ball of radius 2. We obtain a fixed symplectomorphism between an open subset of D^*Q_i and those points in D^*L_i consisting of cotangent vectors lying over a point in L_i of euclidean norm bounded by 4. By identifying the points of D^*Q_1 and D^*Q_2 whose images agree in $D^*L_1 \# D^*L_2$, we obtain a symplectic manifold

$$(2.40) \quad M = D^*Q_1 \#_{(b_1, b_2)} D^*Q_2,$$

which is diffeomorphic to the classical plumbing construction in differential topology [14]. We shall write b for the image of b_1 and b_2 in M , and M_b for the fixed neighborhood of b in M along which the gluing is performed and M_1 and M_2 for the two components of the complement of M_b in M (these are symplectomorphic to the unit disc (cotangent) bundles of the manifolds obtained by removing balls from Q_1 and Q_2).

Since the boundary of M has corners, it is convenient to define a smooth submanifold $M^{\text{in}} \subset M$ as follows: Choose a convex smooth function

$$\chi: [0, +\infty)^2 \rightarrow [0, +\infty)$$

which is a small C^0 perturbation of the maximum of the coordinates and which agrees with this maximum away from a neighborhood of the diagonal. Writing the squared euclidean norms as $\rho(\vec{x}, \vec{y}) = (|\vec{x}|^2, |\vec{y}|^2)$, we consider

$$(2.41) \quad M_0^{\text{in}} = (\chi \circ \rho)^{-1}(\epsilon) \subset D^*L_1 \# D^*L_2;$$

an example for $n = 1$ is shown in Figure 3. The smoothing parameter for ρ should be sufficiently small that the following properties hold:

- 1) Every point in M_0^{in} lies within 2ϵ of a point in $L_1 \cup L_2$.




Figure 3

- 2) If $|\vec{x}| \geq 2\epsilon$, then $(\vec{x}, \vec{y}) \in M_0^{\text{in}}$ if and only if $|\vec{y}| \leq \epsilon$ (and vice versa with x replaced by y).

In words, the second condition says that, away from the disc of radius 2ϵ in L_j , M_0^{in} agrees with the subset of the cotangent bundle consisting of vectors whose length is bounded by ϵ . Writing M_i^{in} for the set of vectors in T^*Q_i of length bounded by ϵ which project to points in Q_i of distance greater than 2ϵ from b_i , we define a smooth manifold with boundary

$$M^{\text{in}} = M_1^{\text{in}} \cup M_b^{\text{in}} \cup M_2^{\text{in}} \subset M.$$

The symplectic form on M can be written as the differential of a primitive θ for which the embedding of each Q_i is exact (i.e., the restriction of θ to Q_i is the differential of a function). Moreover, the tangent space of M is equipped with a homotopy class of complex volume forms (with respect to any almost complex structure compatible with the symplectic form) whose restriction to the cotangent bundle of Q_i is homotopic to a form obtained by complexifying a volume form on Q_i . In particular, Q_i admits a *grading* so that the intersection point between Q_1 and Q_2 is assigned a degree; we omit the choice of grading from our notation, but note that it may be normalized so that

(2.42)

$CF^*(Q_1, Q_2)$ is supported in degree 0.

Definition 2.17. *The category $\text{Fuk}(Q_1, Q_2)$ is the full subcategory of the Fukaya category of M whose objects are Q_1 and Q_2 .*

The existence of an A_∞ category with exact Lagrangians as objects follows from [17]. The construction depends on certain choices of almost complex structures and Hamiltonian perturbations, but the category is independent, up to A_∞ equivalence, of these choices. In order to fix these choices, let us consider a second subdomain M^{mi} of M containing M^{in} , whose construction we shall specify in Section 6 (we shall only use

it to prove compactness for a certain moduli space of holomorphic discs entering in the construction of the functor from Floer to Morse theory). We shall now pick Floer data from the following spaces:

(2.43)

Let \mathcal{J} denote the space of almost complex structures which are of contact type near ∂M^{in} and ∂M^{mi} and let $\mathcal{H} = C_c^\infty(M^{\text{in}}, \mathbb{R})$ denote the space of smooth functions on M whose support is contained in M^{in} .

In the next section we briefly recall how to use perturbations of the $\bar{\partial}$ equation coming from families valued in \mathcal{J} and \mathcal{H} in order to construct $\text{Fuk}(Q_1, Q_2)$.

2.3.1. Review of the construction of the Fukaya category. We follow the discussion in [17], keeping our notation as close as possible. First, we must define the morphism spaces (i.e., Floer complexes).

Definition 2.18. *A Floer datum for each pair (i, j) is the choice of (i) a family $J^{i,j} = \{J_t^{i,j} \in \mathcal{J}\}_{t \in [0,1]}$ of almost complex structures and (ii) a function $H^{i,j} \in \mathcal{H}$.*

This choice determines a perturbed $\bar{\partial}$ equation on the strip

$$(2.44) \quad \bar{\partial}^{i,j} u = (du - Y^{i,j} \otimes dt)^{0,1}$$

where $Y^{i,j}$ is the Hamiltonian flow of $H^{i,j}$, and the $(0, 1)$ part is taken with respect to the t -dependent almost complex structure $J_t^{i,j}$. Given a pair (p_0, p_1) of time-1 chords of the Hamiltonian flow of $H^{i,j}$ with endpoints on Q_j and Q_i , we write

$$(2.45) \quad \begin{aligned} \mathcal{R}(p_0, p_1) = \{u: [0, 1] \times \mathbb{R} \rightarrow M \mid \bar{\partial}^{i,j} u = 0, \\ u(0, s) \in Q_i, u(1, s) \in Q_j, \\ \lim_{s \rightarrow +\infty} u(s, t) = p_0, \lim_{s \rightarrow -\infty} u(s, t) = p_1\} / \mathbb{R} \end{aligned}$$

where the \mathbb{R} action comes from translation on the source. Using the fact that the complex structures which are allowed are of contact type near ∂M^{in} , one can use Lemma 7.2 of [3] (see also Lemma 6.1 below) to prove the analogue of Gromov compactness and conclude:

Lemma 2.19. *If the image of Q_i under the time-1 Hamiltonian flow generated by $H^{i,j}$ is transverse to Q_j , then a generic choice of family $J_t^{i,j}$ ensures that all moduli spaces of gradient trajectories $\mathcal{R}(p_0, p_1)$ for $p_0 \neq p_1$ are regular, and hence of expected dimension*

$$(2.46) \quad \deg(p_0) - \deg(p_1) - 1.$$

Moreover, the Gromov bordification $\bar{\mathcal{R}}(p_0, p_1)$ is compact, and it is a manifold whose boundary, whenever $\deg(p_0) - \deg(p_1) = 2$, is

$$(2.47) \quad \partial \bar{\mathcal{R}}(p_0, p_1) = \coprod_{\deg(p) = \deg(p_1) + 1} \mathcal{R}(p_0, p) \times \mathcal{R}(p, p_1).$$

q. e. d.

The Floer chain complex

$$(2.48) \quad CF^*(Q_i, Q_j)$$

is a direct sum of free abelian groups $|o_p|$ associated to each time-1 chord p from Q_i to Q_j for the Hamiltonian $H^{i,j}$. As the notation suggests, $|o_p|$ is the orientation line associated to a rank 1 space o_p , whose construction we shall not recall. The key fact we use is the existence of a canonical isomorphism

$$(2.49) \quad \lambda(\mathcal{R}(p_0, p_1)) \cong o_{p_0} \otimes o_{p_1}^\vee.$$

In general, the degree $\deg(p)$ of a generator of $CF^*(Q_i, Q_j)$ is determined by an additional choice of grading on Q_i and Q_j . This choice has been fixed (up to a global shift) by requiring that the rank 1 group $CF^*(Q_1, Q_2)$ is supported in degree 0. Whenever $\deg(p_0) - \deg(p_1) = 1$, $\lambda(\mathcal{R}(p_0, p_1))$ is canonically trivialized by the translative action of \mathbb{R} , so Equation (2.49) naturally assigns to each element u of $\mathcal{R}(p_0, p_1)/\mathbb{R}$ a homomorphism

$$\mu_u^F: |o_{p_1}| \rightarrow |o_{p_0}|.$$

The sum of all such homomorphisms is the matrix coefficient of $|o_{p_1}|$ and $|o_{p_0}|$ in the differential in the Floer complex. Writing $[p_1]$ for a generator of $|o_{p_1}|$, we obtain

$$(2.50) \quad \mu_1^F([p_1]) = \sum_{u \in \mathcal{R}(p_0, p_1)/\mathbb{R}} (-1)^{\deg(p_1)} \mu_u^F([p_1]).$$

To define the A_∞ structure, we must choose perturbation data for moduli spaces of discs with an arbitrary number of marked points. We write \mathcal{R}_d for the moduli space of holomorphic discs with $d + 1$ marked points, d of which are marked as incoming, and, given a sequence $\vec{I} = (i_0, \dots, i_d)$, $\mathcal{R}_{\vec{I}}$ for a copy of the moduli space \mathcal{R}_d in which each incoming arc is labeled by an elements of the set $\{1, 2\}$; in particular, the incoming marked points are labeled by pairs (i_k, i_{k+1}) , and the outgoing marked point by (i_0, i_d) . This space admits a compactification $\overline{\mathcal{R}}_{\vec{I}}$ which is a manifold with corners whose points are in bijective correspondence to $\overline{\mathcal{T}}_{\vec{I}}$, and in the setup of Equation (2.26), is equipped with maps

$$(2.51) \quad \overline{\mathcal{R}}_{\vec{I}_1} \times \overline{\mathcal{R}}_{\vec{I}_2} \rightarrow \overline{\mathcal{R}}_{\vec{I}}$$

which cover its boundary. In addition, we choose strip-like ends $\{\epsilon_{i_k}\}_{k=0}^d$ near the marked points of all surfaces Σ in $\overline{\mathcal{R}}_{\vec{I}}$; this choice may be done globally in a smooth family over the moduli space.

Definition 2.20. *A perturbation datum on Σ is a choice of an \mathcal{H} -valued 1-form $K^{\vec{I}}$ on Σ and a function $J^{\vec{I}}: \Sigma \rightarrow \mathcal{J}$. The pull-backs of these data under a strip-like end corresponding to a marked point labeled (i, j) must agree with the Floer data for the pair (i, j) .*

A **universal perturbation datum for discs** $(\mathbf{J}^{\mathcal{R}}, \mathbf{K}^{\mathcal{R}})$ is a choice of perturbation data for all elements of $\overline{\mathcal{R}}_{\vec{I}}$, varying smoothly over the moduli space, and which is consistent, i.e., compatible with the maps (2.51).

The choice of a perturbation datum on Σ defines a perturbed $\bar{\partial}$ equation analogous to (2.44): the Hamiltonian vector field associated to an element of \mathcal{H} and the choice of $K^{\vec{I}}$ define a vector field valued 1-form $Y^{\vec{I}}$ on every surface, so we consider the operator

$$(2.52) \quad \bar{\partial}_{\Sigma}^{\vec{I}}(u) = \left(du - Y^{\vec{I}} \right)^{(0,1)}$$

where the $(0, 1)$ part is taken with respect to the almost complex structure $J^{\vec{I}}$ on the target (and of course the underlying complex structure on Σ on the source).

In particular, given a universal perturbation datum, a sequence $(p_0, \vec{p}) = (p_0, p_1, \dots, p_d)$ with p_0 a chord from Q_{i_0} to Q_{i_d} , and p_k for $1 \leq k \leq d$ chords from $Q_{i_{k-1}}$ to Q_{i_k} (for the appropriate Hamiltonians), we define

$$(2.53) \quad \mathcal{R}(p_0, \vec{p}) = \coprod_{\Sigma \in \overline{\mathcal{R}}_{\vec{I}}} \left\{ u: \Sigma \rightarrow M \mid u(\partial\Sigma) \subset Q_1 \cup Q_2, \right. \\ \left. \lim_{s \rightarrow \pm\infty} u(\epsilon_{i_k}(s, t)) = p_{i_k}, \bar{\partial}_{\Sigma}^{\vec{I}} u = 0 \right\},$$

where the condition on the boundary is more precisely stated as follows: the image of a segment in $\partial\Sigma$ labeled by i_k is contained in Q_{i_k} .

The following result generalizes Lemma 2.19 to discs with multiple inputs and follows from the results proved in [17]:

Proposition 2.21. *For a generic choice of perturbation data, all moduli spaces $\mathcal{R}(p_0, \vec{p})$ are regular and hence have the expected dimension*

$$(2.54) \quad d - 2 + \deg(p_0) - \sum_{1 \leq k \leq d} \deg(p_k).$$

Their Gromov bordifications $\overline{\mathcal{R}}(p_0, \vec{p})$ are compact, and those which have expected dimension 1 are manifolds with boundary. q.e.d.

The gluing theorem for elliptic operators on the disc gives an isomorphism

$$(2.55) \quad \lambda(\mathcal{R}(p_0, \vec{p})) \cong \lambda(\mathcal{R}_d) \otimes \mathfrak{o}_{p_0} \otimes \mathfrak{o}_{p_1}^{\vee} \otimes \dots \otimes \mathfrak{o}_{p_d}^{\vee},$$

where λ is the top exterior power, and hence every rigid holomorphic disc u determines an isomorphism

$$(2.56) \quad \mathfrak{o}_{p_1} \otimes \dots \otimes \mathfrak{o}_{p_d} \rightarrow \mathfrak{o}_{p_0}.$$

Writing μ_u for the induced map on the orientation lines, we use Seidel's conventions, and define the d th higher product

$$\mu_d^F: CF^*(Q_{i_{d-1}}, Q_{i_d}) \otimes \cdots \otimes CF^*(Q_{i_1}, Q_{i_2}) \rightarrow CF^*(Q_{i_1}, Q_{i_d})$$

as a sum of the contributions of all holomorphic discs

$$[p_d] \otimes \cdots \otimes [p_1] \mapsto (-1)^{\dagger(\vec{p})} \sum_{\substack{p_0 \\ u \in \mathcal{R}(p_0, \vec{p})}} \mu_u([p_d] \otimes \cdots \otimes [p_1])$$

where the sign is given by

$$(2.57) \quad \dagger(\vec{p}) = \sum_{k=1}^d k \deg(p_k).$$

3. From simplicial to Morse cochains

In this section, we construct an A_∞ equivalence

$$(3.1) \quad \mathcal{F}: \text{Simp}(Q_1, Q_2) \rightarrow \text{Morse}(Q_1, Q_2).$$

The main ingredients in the construction of such an equivalence is the interpretation of the cup product as an appropriate (perturbed) intersection product, discussed in Section 2.1, and a moduli space of shrubs \mathcal{S}_d , introduced in [2]; it seems that our spaces are combinatorially equivalent to those later introduced by Forcey [7] under the name composihedra. These moduli spaces are quotients of multiplihedra and control A_∞ functors from differential graded algebras to A_∞ algebras. From now on, we shall assume the following

(3.2)

The cells of \check{Q}_i intersect the boundary of \bar{U}_i transversely, and there are nested weak homotopy equivalences $\check{\Delta}_i \subset U_i \subset \Delta_i$. Moreover, the identification $\Delta_1 \cong \Delta_2$ restricts on U_i to the diffeomorphism of Equation (2.20).

In particular for $i \neq j$, we obtain a cellular subdivision $\check{Q}_{i,j}$ of $Q_{i,j}$, which is transverse to the boundary, by pulling back the dual subdivision to Q_i , and the diffeomorphism $Q_{1,2} \cong Q_{2,1}$ respects this cellular subdivision.

We shall briefly review the construction of \mathcal{S}_d and focus on the choices of perturbation data which are necessary to bypass the transversality problems that would arise if we use only gradient flow lines as in [2].

Definition 3.1. *The moduli space of shrubs \mathcal{S}_d is the space of isomorphism classes of metric ribbon trees (S, g_S) with one infinite outgoing edge, and d finite incoming edges whose endpoints are equidistant to the outgoing edge.*

The moduli space of shrubs admits a natural compactification $\bar{\mathcal{S}}_d$ by allowing finite edges to become infinite, or incoming edges to shrink to 0 length. The property of equidistance required of the outgoing

leaves implies that whenever a sequence in \mathcal{S}_d converges to a point on the boundary where the length of some finite edge becomes infinite, every path from the outgoing leaf to an incoming one must contain such an edge. We can reconstruct such limit points as follows: whenever $\sum_{k=1}^r d_k = d$ we obtain a singular shrub by attaching the outputs of r shrubs (the k th shrub having d_k inputs) to a Stasheff tree with r inputs. This construction defines a map

$$(3.3) \quad \overline{\mathcal{T}}_r \times \overline{\mathcal{S}}_{d_1} \times \dots \times \overline{\mathcal{S}}_{d_r} \rightarrow \overline{\mathcal{S}}_d.$$

Similarly, whenever the length of some incoming edge shrinks to 0, there must be at least one other (adjacent) edge which is also collapsed, so that we have for each integer k between 1 and $d - 1$ a map

$$(3.4) \quad \vee_k: \overline{\mathcal{S}}_{d-1} \rightarrow \overline{\mathcal{S}}_d$$

obtained by grafting a pair of edges of length 0 at the k^{th} external vertex. The following result appears in Appendix B of [2]:

Lemma 3.2. *The moduli space $\overline{\mathcal{S}}_d$ is a compact manifold with boundary whose boundary is covered by the images of the maps (3.3) and (3.4).*

As in the previous section, a sequence \vec{I} induces a unique labeling, compatible with the cyclic order, of all edges by pairs $(i(e), j(e))$, such that the incoming edges are labeled by (i_k, i_{k+1}) . Moreover, the maps (3.4) and (3.3) have their analogues for labeled shrubs

$$(3.5) \quad \overline{\mathcal{T}}_{\vec{R}} \times \overline{\mathcal{S}}_{\vec{I}^1} \times \dots \times \overline{\mathcal{S}}_{\vec{I}^r} \rightarrow \overline{\mathcal{S}}_{\vec{I}},$$

$$(3.6) \quad \overline{\mathcal{S}}_{\vec{I}-\{i_k\}} \rightarrow \overline{\mathcal{S}}_{\vec{I}},$$

where the labels on the output of a shrub in $\overline{\mathcal{S}}_{\vec{I}^k}$ agree with the labels of the k th incoming leaf of a Stasheff tree in $\overline{\mathcal{T}}_{\vec{R}}$. A choice of \vec{I} determines a gradient flow equation on every edge of a shrub (S, g_S) , which we perturb as follows:

Definition 3.3. *A universal and consistent perturbation datum for shrubs is a choice \mathbf{X}^S of a smooth family of perturbation data on $\coprod_{\vec{I}} \overline{\mathcal{S}}_{\vec{I}}$ which are compatible with the maps (3.5) and (3.6).*

The discussion of smoothness is relegated to Section 7.1. Recall that the definition of a perturbation datum requires that the resulting perturbed negative gradient flow be outward (respectively, inward) pointing on the boundary of $Q_{1,2}$ (respectively, $Q_{2,1}$). The choice of dual cell subdivision gives another condition:

Definition 3.4. *A universal perturbation datum is compatible with \check{Q} if, whenever e_k and e_{k+1} are successive incoming leaves of length 0 in a shrub (S, g_S) , the vector field*

$$(3.7) \quad -\nabla f_{i_k, i_{k+1}} + X_{e_{k+1}} - (-\nabla f_{i_{k-1}, i_k} + X_{e_k})$$

defined on the intersection of the images of Q_{i_{k-1}, i_k} and $Q_{i_k, i_{k+1}}$ in Q is simplicial.

By Condition (3.2) the behavior of the (perturbed) gradient flow near the boundary is compatible with a requirement that the vector field be simplicial. Indeed, the restriction of the tangent space of $Q_{i,j}$ to every cell $\tilde{\tau}$ of $\tilde{Q}_{i,j}$ contains an open cone C_τ consisting of vectors pointing in the direction of the top-dimensional cell adjacent to $\tilde{\tau}$ which is minimal with respect to the ordering; the simplicial condition is the requirement that the restriction of a vector field to $\tilde{\tau}$ lie in C_τ . As the boundary of C_τ is covered by the tangent spaces of the codimension 1 faces meeting at $\tilde{\tau}$, $T\tilde{\tau}$ is included in its closure. Since $\partial Q_{i,j}$ intersects $\tilde{\tau}$ transversely, we conclude that the tangent space of the boundary separates C_τ into two non-empty open components. Assuming that $f_{i_k, i_{k+1}}$, f_{i_{k-1}, i_k} , and X_{e_k} have been chosen already, we can always choose $X_{e_{k+1}}$ deep inside the appropriate cone so that (3.7) is a simplicial vector field.

Lemma 3.5. *There is a non-empty open subset of the space of universal consistent perturbation data for shrubs which consists of data that are compatible with \tilde{Q} .*

Proof. Since universal perturbation data are constructed inductively, it suffices to see that the compatibility condition on $\overline{\mathcal{S}}_{\tilde{I}}$ does not restrict the choice of perturbation data for $\overline{\mathcal{S}}_{\tilde{I}'}$ whenever \tilde{I}' is a subset of \tilde{I} . This is essentially obvious since our perturbation data on each edge are chosen independently of each other and the compatibility condition restricts edges in $\overline{\mathcal{S}}_{\tilde{I}}$ which do not exist in $\overline{\mathcal{S}}_{\tilde{I}'}$. q.e.d.

From now on when we speak of a universal perturbation datum for shrubs, we shall assume that it is compatible with \tilde{Q} . With the choice of such a datum, we can now define the moduli spaces which shall be used to construct the functor of Equation (3.1). Let us write $\vec{\sigma} = (\check{\sigma}_1, \dots, \check{\sigma}_d)$ for a sequence of cells $\check{\sigma}_k \in \tilde{Q}_{i_k, i_{k+1}}$, and x_0 for a critical point of f_{i_0, i_d} :

Definition 3.6. *The moduli space $\mathcal{S}(x_0, \vec{\sigma})$ of (perturbed) gradient shrubs with inputs $\vec{\sigma}$ and output x_0 is the set of maps*

$$(3.8) \quad \psi: (S, g_S) \rightarrow Q$$

such that (i) each edge e of S lifts as a perturbed gradient flow segment to the labeling manifold $Q_{i(e), j(e)}$, (ii) the lift of the outgoing leaf is the critical point x , and (iii) the lift of the k th incoming leaf lies on $\check{\sigma}_k$.

The moduli space $\mathcal{S}(x_0, \vec{\sigma})$ admits a natural stratification coming from the stratification of \mathcal{S}_d by the topological type of the tree, and a bordification $\overline{\mathcal{S}}(x, \vec{\sigma})$ by allowing edges to become infinite. In particular, we have a subset

$$(3.9) \quad \overline{\mathcal{S}}_{v_k}(x_0, \vec{\sigma}) \subset \overline{\mathcal{S}}(x_0, \vec{\sigma})$$

consisting of maps whose domain lies in the image of the inclusion of Equation (3.6); i.e. singular shrubs whose k and $(k + 1)$ st incoming edges both have length 0.

Lemma 3.7. *If the cup product of $\check{\sigma}_k$ and $\check{\sigma}_{k+1}$ vanishes, then there is a neighborhood of the image of \vee_k in $\text{int}(\overline{\mathcal{S}}_{\vec{\sigma}})$ whose inverse image under the forgetful map*

$$(3.10) \quad \overline{\mathcal{S}}(x_0, \vec{\sigma}) \rightarrow \overline{\mathcal{S}}_{\vec{\sigma}}$$

is empty. In particular, this stratum of $\overline{\mathcal{S}}(x_0, \vec{\sigma})$ does not lie in the closure of the top-dimensional part of the moduli space $\mathcal{S}(x_0, \vec{\sigma})$.

Proof. The proof is completely local. A slight generalization of Lemma 2.3 implies that the condition that the cup product of $\check{\sigma}_k$ and $\check{\sigma}_{k+1}$ vanish is equivalent to the absence of intersection points between the images of $\check{\sigma}_k$ under a flow ϕ_k and $\check{\sigma}_{k+1}$ under a flow ϕ_{k+1} , whenever ϕ_k and ϕ_{k+1} are generated by sufficiently small vector fields X_k and X_{k+1} whose difference is compatible with the simplicial triangulation. If (S, g_S) is a singular shrub with sufficiently short incoming edges e_k and e_{k+1} , the flows obtained by integrating the perturbed gradient vector field along these edges satisfy this property due to our assumption that the perturbation datum $X^{\mathcal{S}}$ is smooth and compatible with the subdivision \tilde{Q} . The desired result follows immediately. q.e.d.

We also need a compactness result analogous to Lemma 2.12:

Lemma 3.8. *Given x_0 a critical point of $f_{0,d}$, and $(\check{\sigma}_1, \dots, \check{\sigma}_d)$ cells of the subdivisions $Q_{i_k, i_{k+1}}$, the moduli space $\overline{\mathcal{S}}(x_0, \vec{\sigma})$ is compact.*

Proof. Again, it suffices to prove that no edge labeled by $Q_{1,2}$ or $Q_{2,1}$ may escape to the boundary, which can be done in three cases, of which we explain only one. We fix a neighborhood of the boundary of $Q_{2,1}$ which does not intersect the cell dual to the interior of Δ . Assuming that the image of an edge e labeled by $Q_{2,1}$ intersects this neighborhood, we follow an ascending arc all of whose edges are labeled by $Q_{2,1}$ and derive the contradictory conclusion that an input cell labeled by $(1, 2)$ intersects a neighborhood of the boundary of $Q_{2,1}$. q.e.d.

In Section 7.1, we explain how to prove:

Lemma 3.9. *$\overline{\mathcal{S}}(x, \vec{\sigma})$ is compact, and for generic universal perturbation data, all strata are smooth manifolds of the expected dimension. Moreover, the closure of the top-dimensional strata is a compact manifold with boundary of dimension*

$$(3.11) \quad d - 1 + \deg(x_0) - \sum_k \deg(\check{\sigma}_k),$$

which we denote $\hat{\mathcal{S}}(x_0, \vec{\sigma})$.

We now restrict attention to the case where the moduli space of shrubs is rigid, i.e., to a sequence of cells $\vec{\sigma}$ and a critical point x_0 such that

$$(3.12) \quad \deg(x_0) = 1 - d + \sum_k \deg(\check{\sigma}_k).$$

The description of the moduli space of shrubs as a fibered product (see Section 8.3) gives an isomorphism

$$(3.13) \quad \lambda(\mathcal{S}_d(x_0, \vec{\sigma})) \otimes \lambda(Q^{d+1}) \cong \lambda(Q) \otimes \lambda(\mathcal{S}_d) \otimes \lambda(W^s(x_0)) \otimes \lambda(\vec{\sigma}),$$

where the last factor on the right hand side is a tensor product

$$(3.14) \quad \lambda(\vec{\sigma}) \cong \lambda(\check{\sigma}_1) \otimes \cdots \otimes \lambda(\check{\sigma}_d).$$

Using the isomorphism

$$(3.15) \quad \lambda(Q) \cong \lambda(\sigma_k) \otimes \lambda(\check{\sigma}_k),$$

we conclude that a rigid shrub $\psi \in \mathcal{S}_d(x_0, \vec{\sigma})$ determines an isomorphism

$$(3.16) \quad \lambda(\sigma_1) \otimes \cdots \otimes \lambda(\sigma_d) \rightarrow \mathfrak{o}_{x_0}.$$

Writing \mathcal{F}^ψ for the induced map on the orientation line, the d th order map in our functor from the simplicial to the Morse category is defined via the formula

$$(3.17) \quad \mathcal{F}^d: C^*(\mathcal{Q}_{i_{d-1}, i_d}) \otimes \cdots \otimes C^*(\mathcal{Q}_{i_0, i_1}) \rightarrow CM^*(f_{0,d})$$

$$(3.18) \quad \check{\sigma}_d \otimes \cdots \otimes \check{\sigma}_1 \mapsto (-1)^{(n+1)\dagger(\vec{\sigma})} \sum_{\substack{x_0 \\ \psi \in \mathcal{S}_d(x_0, \vec{\sigma})}} \mathcal{F}^\psi(\vec{\sigma}),$$

where the sign $\dagger(\vec{\sigma})$ is determined by a choice of orientation on the moduli of shrubs. With our sign conventions, we would find

$$(3.19) \quad \dagger(\vec{\sigma}) = \sum_{k=1}^d k \deg(\check{\sigma}_k).$$

To illustrate the case of one input, let us fix a simplex σ_1 of $\mathcal{Q}_{i,j}$ and a critical point x_0 of $f_{i,j}$ such that $\check{\sigma}_1$ and $W^s(x_0)$ intersect transversely. At every such intersection point ψ , we have an isomorphism

$$(3.20) \quad \lambda(Q) \cong \lambda(W^s(x_0)) \otimes \lambda(\check{\sigma}_1),$$

which, upon using the decomposition (3.15) yields an isomorphism

$$(3.21) \quad \lambda(\sigma_1) \cong \lambda(W^s(x_0)).$$

The sign from Equation (3.19) is equal to

$$(3.22) \quad (-1)^{(n+1)\deg(\check{\sigma})}$$

in this case, which determines the contribution of ψ to the map

$$(3.23) \quad \mathcal{F}^1: C^*(\mathcal{Q}_{i,j}) \rightarrow CM^*(f_{i,j}).$$

To prove that this is a chain map, we must consider, as usual, 1-dimensional moduli spaces of shrubs with one input. This occurs whenever we are given a critical point x_0 such that

$$(3.24) \quad \deg(x_0) = \deg(\check{\sigma}_1) + 1.$$

Since the space $\overline{\mathcal{S}}(x_0, \check{\sigma}_1)$ is the closure of the intersection of $\check{\sigma}_1$ with $W^s(x_0)$, its boundary is easily seen to consist of two strata: either (i) a family of gradient trajectories from x_0 escapes to a cell $\check{\sigma}_0$ on the boundary of $\check{\sigma}_1$, or (ii) a family of such trajectories converges to a broken one, corresponding to the concatenation of an ascending gradient flow line from x_0 to a critical point x_1 of degree equal to $\deg(\check{\sigma}_1)$, followed by a gradient flow line from x_1 to $\check{\sigma}_1$. These two configurations correspond to the two terms whose equality implies that \mathcal{F}^1 is a chain map:

$$(3.25) \quad \mathcal{F}^1 \circ \mu_1^S = \mu_1^M \circ \mathcal{F}^1.$$

To prove the correctness of the signs, we note that we have a natural isomorphism

$$(3.26) \quad \lambda(\mathcal{S}(x_0, \check{\sigma}_1)) \otimes \lambda(Q) \cong \lambda(W^s(x_0)) \otimes \lambda(\check{\sigma}_1)$$

which we must compare with the isomorphisms (3.20) and (8.2) when we pass to the compactification $\overline{\mathcal{S}}(x_0, \check{\sigma}_1)$.

First, we obtain an isomorphism

$$(3.27) \quad \lambda(\mathbb{R}) \otimes \lambda(W^s(x_1)) \cong \lambda(W^s(x_0))$$

from Equations (8.2) and (8.4). The result of taking the tensor product with the two sides of Equation (3.20) gives

$$(3.28) \quad \lambda(\mathbb{R}) \otimes \lambda(W^s(x_1)) \otimes \lambda(Q) \cong \lambda(W^s(x_0)) \otimes \lambda(W^s(x_1)) \otimes \lambda(\check{\sigma}_1).$$

The reader may easily check that the gluing theorem for gradient flow lines gives an isomorphism $\lambda(\mathbb{R}) \cong \lambda(\mathcal{S}(x_0, \check{\sigma}_1))$ near a broken trajectory, with the positive direction in \mathbb{R} corresponding to the *outward* pointing normal vector. In order to pass from this to the isomorphism (3.26), we must switch some factors, taking care to record the appropriate Koszul signs:

- 1) Permute the first copy of $\lambda(W^s(x_1))$ past $\lambda(\mathbb{R})$, introducing a parity change of $\deg(\check{\sigma}_1)$.
- 2) Permute the two factors of $\lambda(W^s(x_0)) \otimes \lambda(W^s(x_1))$, introducing a sign equal to $\deg(x_0) \deg(x_1)$, which vanishes modulo 2.
- 3) Cancel the $\lambda(W^s(x_1))$ factors on either side.

The result of the operations above is exactly the isomorphism of Equation (3.26), with a change in signs given by the parity of $\deg(\check{\sigma}_1)$.

We must now analyze the same problem near the other type of boundary strata. Near such a stratum, we have a natural isomorphism between

$\lambda(\mathcal{S}(x_0, \check{\sigma}_1))$ and the normal vector of the boundary to $\check{\sigma}_1$. In particular, Equation (3.26) gives an isomorphism

$$(3.29) \quad \lambda(\mathcal{S}(x_0, \check{\sigma}_1)) \otimes \lambda(Q) \cong \lambda(W^s(x_0)) \otimes \lambda(\mathcal{S}(x_0, \check{\sigma}_1)) \otimes \lambda(\partial\check{\sigma}_1)$$

which, after introducing a sign with parity $\deg(x_0) = \deg(\check{\sigma}_1) + 1$, gives exactly the isomorphism of (3.20) for the boundary of σ . Note that this would naively seem to imply that \mathcal{F}^1 satisfies the desired equation, since the two types of boundary strata have *opposite relative orientations*.

Unfortunately, there are additional signs coming from (i) the signs in formulas (8.7) and (2.16) for the differentials in the Morse and simplicial categories and (ii) the sign in the definition of the map \mathcal{F}^1 . Remark 2.2 explains that the sign introduced in the definition of $\text{Simp}(\mathcal{Q}_1, \mathcal{Q}_2)$ contributes a -1 to the total formula because it corrects for the fact that dual cells do not have compatible orientations on their boundaries. On the other hand, both the Morse differential and \mathcal{F}^1 contribute $(-1)^n$ and $(-1)^{n+1}$, respectively, so the artificial signs we have introduced cancel in this case. The generalization for higher products, gives:

Proposition 3.10. *The maps of Equation (3.17) define an A_∞ equivalence. In particular, \mathcal{F}^1 is a chain map inducing an isomorphism on cohomology, and*

$$(3.30) \quad \sum_{1 \leq k \leq d-1} (-1)^{\mathfrak{X}_k} \mathcal{F}^{d-1} \left(\text{id}^{d-k-2} \otimes \mu_2^S \otimes \text{id}^k \right) + \sum_{1 \leq k \leq d} (-1)^{\mathfrak{X}_k} \mathcal{F}^d \left(\text{id}^{d-k-1} \otimes \mu_1^S \otimes \text{id}^k \right) = \sum_{d_1 + \dots + d_r = d} \mu_r^M \left(\mathcal{F}^{d_1} \otimes \dots \otimes \mathcal{F}^{d_r} \right)$$

where

$$\mathfrak{X}_k(\vec{p}) = k + \sum_{j=1}^k \deg(p_j).$$

Proof. The claim that \mathcal{F}^1 is a chain isomorphism is the familiar statement that the cohomology of smooth manifolds can be computed using Morse theory. To prove the A_∞ equation (3.30), we study, as usual, moduli spaces $\mathcal{S}(x, \vec{\sigma})$ of expected dimension 1. By Lemma 3.9 the closure of the interior of this moduli space is a 1-dimensional manifold with boundary projecting to the interior or top-dimensional strata of the boundary of $\overline{\mathcal{S}}_d$ (this follows from the fact that all strata of $\mathcal{S}(x, \vec{\sigma})$ are of the expected dimension, and hence that those lying over higher codimension strata of $\overline{\mathcal{S}}_d$ are empty).

It suffices to check that the boundary points of $\hat{\mathcal{S}}(x, \vec{\sigma})$ account for all the terms of (3.30). The first sum in the left-hand side of (3.30) corresponds to boundary points of $\hat{\mathcal{S}}(x, \vec{\sigma})$ which project to the interior of $\overline{\mathcal{S}}_d$. The top-dimensional boundary strata of $\overline{\mathcal{S}}_d$ are of two types: those which are the image of (3.3) correspond to the right-hand side of

(3.30), and we claim that the second term of the left-hand side of that equation correspond to those points which project to the image of (3.4) and are in the closure of $\mathcal{S}(x, \vec{\sigma})$. Indeed, points in $\overline{\mathcal{S}}(x, \vec{\sigma})$ lying over the image of (3.4) correspond to $\mathcal{S}(x, \vec{\sigma}^{\cap k})$, where $\vec{\sigma}^{\cap k}$ is obtained from $\vec{\sigma}$ by replacing the cells $\check{\sigma}_k$ and $\check{\sigma}_{k+1}$ by their *intersection*. Lemma 3.7 implies that whenever this intersection does not contribute to the cup product, the corresponding cell does not lie in the boundary of $\hat{\mathcal{S}}(x, \vec{\sigma})$, so, modulo signs, the Proposition follows from the next Lemma. We briefly discuss signs and orientations in Section 8.3. q.e.d.

Lemma 3.11. *If $\mu_2^C(\check{\sigma}_k, \check{\sigma}_{k+1}) \neq 0$, and the moduli space $\mathcal{S}(x, \vec{\sigma}^{\cap k})$ is non-empty, the corresponding points of $\overline{\mathcal{S}}(x, \vec{\sigma})$ lie in the boundary of $\hat{\mathcal{S}}(x, \vec{\sigma})$.*

Proof. For the purpose of this argument we introduce a moduli space

$$(3.31) \quad \mathcal{S}_{d-1}^{k, \epsilon}$$

consisting of ribbon trees satisfying all the conditions for being a shrub except that the distance from the outgoing edge to the incoming leaves numbered k and $k+1$ are smaller by ϵ than the distances to the other incoming vertices. Note that the regularity of $\mathcal{S}(x, \vec{\sigma}^{\cap k})$ and the implicit function theorem imply that, whenever ϵ is sufficiently small, we have a bijective correspondence between its elements and those of $\mathcal{S}^{k, \epsilon}(x, \vec{\sigma}_\phi^{\cap k})$, where $\vec{\sigma}_\phi^{\cap k}$ is obtained from $\vec{\sigma}^{\cap k}$ by replacing $\check{\sigma}_k \cap \check{\sigma}_{k+1}$ by its image under a diffeomorphism sufficiently close (in the C^1 norm) to the identity, and keeping all other cells unchanged. Using the fact that our chosen perturbations are compatible with the subdivision \tilde{Q} , we find that whenever $\mu_2^C(\check{\sigma}_k, \check{\sigma}_{k+1}) \neq 0$, there is a diffeomorphism $\phi_{k, k+1}^{S, g_S}$ such that

$$(3.32) \quad \phi_{k, k+1}^{S, g_S}(\check{\sigma}_k \cap \check{\sigma}_{k+1}) = \phi_k^{S, g_S}(\check{\sigma}_k) \cap \phi_{k+1}^{S, g_S}(\check{\sigma}_{k+1})$$

where ϕ_k^{S, g_S} and ϕ_{k+1}^{S, g_S} are the diffeomorphisms obtained by integrating the perturbed gradient flow along the respective incoming leaves. Moreover, as (S, g_S) converges to the relevant boundary stratum of $\overline{\mathcal{S}}_d$, the map $\phi_{k, k+1}^{S, g_S}$ may be chosen to converge to the identity. Writing $\epsilon_k(S, g_S)$ for the length of the k^{th} incoming leaf, the claim follows by noticing that the elements of

$$\mathcal{S}^{k, \epsilon_k(S, g_S)} \left(x, \vec{\sigma}_{\phi_{k, k+1}^{S, g_S}}^{\cap k} \right)$$

are in bijective correspondence with the elements of the top-dimensional stratum of $\mathcal{S}_{(S, g_S)}(x, \vec{\sigma})$ which are sufficiently close to $\mathcal{S}(x, \vec{\sigma}^{\cap k})$. q.e.d.




Figure 4. The unique mushroom with one input, with cap P and stem S .

4. A convenient model for Multiplihedra

Boardman and Vogt [5] introduced spaces of “colored trees” as a model for multiplihedra; this collections of polyhedra controls A_∞ functors. Since our goal in later sections is to construct an A_∞ functor whose source is the Fukaya category, it seems reasonable to replace one of the trees by holomorphic discs. Keeping relatively close to our previous botanical terminology, we shall call these spaces moduli spaces of mushrooms.

Definition 4.1. A **mushroom** C with d inputs is defined by the following data:

- 1) A partition $d = d_1 + d_2 + \cdots + d_r$ (in particular $d_k \geq 1$)
- 2) For each integer $1 \leq k \leq r$, a **cap** $P_k(C)$ which is a disc with $d_k + 2$ marked boundary points of which d_k successive ones are distinguished as incoming
- 3) A **stem** $S(C)$, which is a a shrub with k inputs

REMARK 4.2. Sometimes, it shall be convenient to discuss caps separately. By a d -cap, we will simply mean a disc with $d + 2$ marked points of which d successive ones are distinguished as incoming. Alternatively, we can distinguish the pair of outgoing vertices, or the arc connecting them. The **moduli space of caps** \mathcal{P}_d is simply the (open) Stasheff moduli space \mathcal{T}_{d+1} under another guise. Later, we shall find it convenient to fix the identification with the moduli space of discs \mathcal{R}_{d+1} for which the boundary marked point v_1^{out} in Figure 4 is mapped to the (unique) outgoing marked point for discs in \mathcal{R}_{d+1} .

REMARK 4.3. Even though we will draw mushrooms as in Figure 4, we remind the reader that the incoming vertices of the stem are not associated to a marked point on the outgoing arc.

An isomorphism $C \rightarrow C'$ of mushrooms is an isomorphism of the underlying caps and stem. Explicitly, such a map is a collection of $r + 1$

isomorphisms:

$$\begin{aligned} P_k(C) &\rightarrow P_k(C') \quad \text{for } 1 \leq k \leq r \\ S(C) &\rightarrow S(C'). \end{aligned}$$

Lemma 4.4. *The only non-trivial automorphism of a mushroom is the identity.*

Proof. Such an automorphism would correspond to an automorphism of the stem or of one of the caps. However, we have excluded discs with less than 3 marked points as caps. q.e.d.

By definition, the set \mathcal{C} of mushrooms with d inputs admits a bijective map to

$$(4.1) \quad \coprod \mathcal{S}_r \times \mathcal{P}_{d_1} \times \cdots \times \mathcal{P}_{d_r}$$

which we later use to define a topology. Recall that whereas the moduli space \mathcal{T}_d is an open manifold, the moduli space of shrubs is a (non-compact) manifold with boundary (the boundary consists of trees with outgoing edges of vanishing length). Note that the dimension of the components of (4.1) depends only on the number of inputs:

$$(r - 1) + \sum_{i=1}^r (d_i - 1) = d - 1.$$

Definition 4.5. *A singular mushroom with d inputs is defined by the following data:*

- 1) *A partition $d = d_1 + d_2 + \cdots + d_r$ with $d_k \geq 1$*
- 2) *For each integer $1 \leq k \leq r$, a **cap**, which is a singular disc with $d_k + 2$ marked boundary points with a distinguished arc connecting two successive (outgoing) marked points*
- 3) *A **stem**, which is a singular shrub with r inputs*

REMARK 4.6. Again, we consider singular caps whose moduli spaces $\overline{\mathcal{P}}_d$ can be identified with the Stasheff polytopes in their holomorphic disc model $\overline{\mathcal{R}}_{d+1}$. Decomposing a singular cap into its irreducible components, we observe that those components whose boundary intersects the distinguished arc can themselves be interpreted as lower-dimensional caps. Not all irreducible components are of this type (see Figure 5).

As usual, we will focus attention on **stable** singular mushrooms. These are characterized by stable caps and stem. As before, we can write this space as

$$\coprod \overline{\mathcal{S}}_r \times \overline{\mathcal{P}}_{d_1} \times \cdots \times \overline{\mathcal{P}}_{d_r}.$$

However, it is natural to think of the moduli space of stable singular mushrooms as a quotient of this space.




Figure 5. A singular cap consisting of two caps and an “ordinary” disc.

For each d' and d'' , we consider the embedding \circ

$$(4.2) \quad \overline{\mathcal{P}}_{d'} \times \overline{\mathcal{P}}_{d''} \rightarrow \partial \overline{\mathcal{P}}_{d'+d''} \hookrightarrow \overline{\mathcal{P}}_{d'+d''}$$

whose image is the stratum of $\partial \overline{\mathcal{P}}_{d'+d''}$ consisting of discs which can be decomposed into two components (not necessarily irreducible), each of which contains an outgoing marked point, and which have, respectively, d' and d'' incoming marked points.

On the other hand, for each integer $1 \leq k \leq r-1$, we have a map

$$\vee_k: \overline{\mathcal{S}}_{r-1} \rightarrow \overline{\mathcal{S}}_r$$

which creates two incoming edges of vanishing length at the k^{th} vertex. We therefore have maps

$$\begin{array}{c} \overline{\mathcal{S}}_r \times \overline{\mathcal{P}}_{d_1} \times \cdots \times \overline{\mathcal{P}}_{d_{k-1}} \times \overline{\mathcal{P}}_{d_k} \times \cdots \times \overline{\mathcal{P}}_{d_r} \\ \uparrow \vee_k \times \text{id}^r \\ \overline{\mathcal{S}}_{r-1} \times \overline{\mathcal{P}}_{d_1} \times \cdots \times \overline{\mathcal{P}}_{d_{k-1}} \times \overline{\mathcal{P}}_{d_k} \times \cdots \times \overline{\mathcal{P}}_{d_r} \\ \downarrow \text{id}^{k-1} \times \circ \times \text{id}^{r-k} \\ \overline{\mathcal{S}}_{r-1} \times \overline{\mathcal{P}}_{d_1} \times \cdots \times \overline{\mathcal{P}}_{d_{k-1}+d_k} \times \cdots \times \overline{\mathcal{P}}_{d_r}. \end{array}$$

We say that two stable singular mushrooms are **equivalent** if they are identified upon gluing along the maps $\vee_k \times \text{id}$ and $\text{id} \times \circ \times \text{id}$. We write \sim for the equivalence relation generated by all such diagrams:

Definition 4.7. *The moduli space of stable mushrooms with d inputs is the quotient space*

$$\overline{\mathcal{C}}_d = \coprod \overline{\mathcal{S}}_r \times \overline{\mathcal{P}}_{d_1} \times \cdots \times \overline{\mathcal{P}}_{d_r} / \sim.$$

We will equip \mathcal{C}_d with its inherited topology as a subset of $\overline{\mathcal{C}}_d$. The spaces $\overline{\mathcal{C}}_d$ admit two important families of maps of operadic type relating them to the other moduli spaces we have studied. On the one hand, by grafting r mushrooms at the endpoints of a Stasheff tree, we obtain a map

$$(4.3) \quad \overline{\mathcal{T}}_r \times \overline{\mathcal{C}}_{d_1} \times \cdots \times \overline{\mathcal{C}}_{d_r} \rightarrow \overline{\mathcal{C}}_d$$




Figure 6. Singular mushrooms lying on two different boundary strata of $\overline{\mathcal{C}}_4$.

where $\sum d_j = d$. Alternatively, by attaching a singular disc at any of the incoming marked point, we obtain d_1 maps

$$(4.4) \quad \overline{\mathcal{C}}_{d_1} \times \overline{\mathcal{R}}_{d_2} \rightarrow \overline{\mathcal{C}}_{d_1+d_2-1}.$$

The moduli space $\overline{\mathcal{R}}_{d+1}$ is stratified by cells labeled by topological types of ribbon trees with $d + 2$ leaves, which induces a stratification of $\overline{\mathcal{P}}_d$. When considering this latter space, we shall assume that two such successive leaves are distinguished as outgoing. Note that the trees in which the two outgoing leaves are not adjacent (i.e., share a vertex) are exactly those that label singular caps whose outgoing edge meets more than one irreducible component. The reader may compare this constraint with the condition imposed by Boardman and Vogt on their colored trees [5] (see also, [7]), which shows that our moduli space $\overline{\mathcal{C}}_d$ is homeomorphic to the multiplihedron of the same dimension. Moreover, the homeomorphism may be chosen so that the maps (4.3) and (4.4) are the operadic structure maps which make multiplihedra into a bimodule over the Stasheff operad. We conclude:

Lemma 4.8. *$\overline{\mathcal{C}}_d$ is a compact manifold with boundary stratified into finitely many (smooth) manifolds with corners. The top boundary strata are the images of the map (4.3) and (4.4) (see Figure 6). q.e.d.*

Note that the top boundary strata of $\overline{\mathcal{C}}_d$ are the boundary strata of

$$\coprod \overline{\mathcal{S}}_r \times \overline{\mathcal{P}}_{d_1} \times \cdots \times \overline{\mathcal{P}}_{d_r}$$

which are not identified by the equivalence relation. In fact, the top boundary strata which are identified by the equivalence relation are exactly those for which a pair of successive external edges have vanishing length. It is easy to check that the remaining strata are exactly as described above.

As with trees (and discs), this construction can be done with labels. Given a sequence $\vec{I} = (i_0, \dots, i_d)$, we shall consider the moduli space of mushrooms labeled by \vec{I} , denoted $\bar{\mathcal{C}}_{\vec{I}}$, which is a copy of $\bar{\mathcal{C}}_d$. We start by setting $\bar{\mathcal{P}}_{\vec{I}^k}$ to be a copy of $\bar{\mathcal{P}}_{|\vec{I}^k|-1}$ where the incoming arcs are labeled by the sequence \vec{I}^k , which allows us to define $\bar{\mathcal{C}}_{\vec{I}}$ to be the quotient (under the identification of used in Definition 4.7) of the disjoint union

$$(4.5) \quad \coprod \bar{\mathcal{S}}_{\vec{R}} \times \bar{\mathcal{P}}_{\vec{I}^1} \times \cdots \times \bar{\mathcal{P}}_{\vec{I}^r}$$

where the initial elements of \vec{I}^k agrees with the last element of \vec{I}^{k-1} , the sequence \vec{I} is obtained by concatenating \vec{I}^1 through \vec{I}^r and dropping the redundant terms, and \vec{R} is the sequence of length $r+1$ obtained by taking i_0 together with the last elements of \vec{I}^1 through \vec{I}^r . Whenever \vec{R} and $(\vec{I}^1, \dots, \vec{I}^r)$ are any sequences satisfying this condition, we obtain a map

$$(4.6) \quad \bar{\mathcal{T}}_{\vec{R}} \times \bar{\mathcal{C}}_{\vec{I}^1} \times \cdots \times \bar{\mathcal{C}}_{\vec{I}^r} \rightarrow \bar{\mathcal{C}}_{\vec{I}}$$

which is the analogue of (4.3) in the presence of labels. The analogue of Equation (4.4) is a map

$$(4.7) \quad \bar{\mathcal{C}}_{\vec{I}^1} \times \bar{\mathcal{R}}_{\vec{I}^2} \rightarrow \bar{\mathcal{C}}_{\vec{I}}$$

where \vec{I} is obtained from \vec{I}^1 and \vec{I}^2 as in Equation (2.26). This later map is induced by a map

$$(4.8) \quad \bar{\mathcal{P}}_{\vec{I}^1} \times \bar{\mathcal{R}}_{\vec{I}^2} \rightarrow \bar{\mathcal{P}}_{\vec{I}}.$$

Finally, the gluing operation (4.2) also has a labeled analogue

$$(4.9) \quad \bar{\mathcal{P}}_{\vec{I}^1} \times \bar{\mathcal{P}}_{\vec{I}^2} \rightarrow \bar{\mathcal{P}}_{\vec{I}}$$

where the last element of \vec{I}^1 is assumed to agree with the first element of \vec{I}^2 , and \vec{I} is obtained by concatenating the two sequences and dropping the repeated term.

5. From Floer to Morse cochains

In this section, we construct a functor

$$(5.1) \quad \mathcal{G}: \text{Fuk}(Q_1, Q_2) \rightarrow \text{Morse}(Q_1, Q_2)$$

using “moduli spaces of mushrooms maps,” which we shall define in Section 5.3, and a Lagrangian foliation on the plumbing, parametrized by Q , which we shall define in Section 5.1 (recall that Q is the space obtained by gluing Q_1 and Q_2 along open balls). Informally, a mushroom map consists of a gradient tree whose source is the stem, and a holomorphic map with Lagrangian boundary conditions on each cap. The Lagrangian condition on the boundary components connecting incoming vertices (or an incoming vertex and an outgoing one) will be

one of the Lagrangians Q_i that we have been studying. The outgoing boundary component of a cap (the one connecting the two outgoing vertices) will be required to map to a leaf of the Lagrangian foliation on M whose label (a point in Q) agrees with the image of the corresponding endpoint of the stem. From now on, we shall assume:

(5.2)

The open sets U_i used to define the Morse category are the images of the balls of radius 2 in \mathbb{R}^n under the charts used in Section 2.3.

5.1. A Lagrangian foliation. Returning to the local model of Section 2.3, we find that the submanifold

$$(5.3) \quad L_\Delta = \{x_i = y_i\}_{i=1}^n$$

is also Lagrangian, and given $\vec{x} \in \mathbb{R}^n$, we let $(\vec{x}, 0) + L_\Delta$ denote the translate by $(\vec{x}, 0)$ of L_Δ ; the union of all such leaves is a foliation of \mathbb{C}^n . Note that a fixed leaf $(\vec{x}, 0) + L_\Delta$ intersects L_2 transversely at $(0, -\vec{x})$. Using Stokes's theorem to integrate the action of a path from L_1 to L_2 along $(\vec{x}, 0) + L_\Delta$, we conclude:

Lemma 5.1. *The integral of ω over any disc with counterclockwise boundary conditions $(L_2, L_1, (\vec{x}, 0) + L_\Delta)$ is proportional to $-|x|^2$; in particular, no such disc is holomorphic if $x \neq 0$, in which case only constant discs can arise. q.e.d.*

A minor deficiency with this model is that it differs from the standard fibration of \mathbb{C}^n as the cotangent bundle of L_1 or L_2 , so we must modify it in order to extend it to M . We write L for the result of gluing L_1 to L_2 along the identification of their balls of radius 1 about the origin under the diffeomorphism

$$(5.4) \quad (\vec{x}, 0) \mapsto (0, -\vec{x}).$$

This space is a local model for Q in the sense that there is an open subset of Q , containing all non-Hausdorff points, which is diffeomorphic to L , and which maps L_1 to the image of Q_1 and L_2 to the image of Q_2 .

We would like to define a foliation of $D^*L_1 \# D^*L_2$, whose leaves \mathcal{L}_ℓ are Lagrangian discs parametrized by $\ell \in L$. and such that the following essential properties hold:

- \mathcal{L}_ℓ intersects L_1 at $(\vec{x}, 0)$ if and only if $(\vec{x}, 0)$ represents ℓ in L (with the same property for L_2).
- If $4 < |\vec{x}|$, then $\mathcal{L}_{(\vec{x}, 0)}$ and $\mathcal{L}_{(0, \vec{x})}$ are respectively given by the affine Lagrangians orthogonal to L_1 and L_2 at $(\vec{x}, 0)$ and $(0, \vec{x})$.

The first condition gives the correspondence between leaves and points of L , while the second implies that this foliation agrees with the canonical foliation of the cotangent bundles of L_1 and L_2 away from a neighborhood of the origin, and hence can be extended to M . In addition, we




Figure 7

require three technical conditions which will be used to prove compactness for moduli spaces of holomorphic curves (see Section 6), and which depend on parameters $\epsilon \ll \delta$ which can be made arbitrarily small; ϵ already appeared in the construction of M^{in} Section 2.3.1:

(5.5)

If $\delta \leq |\vec{x}|$ then the intersections $\mathcal{L}_{(\vec{x},0)} \cap M_1^{\text{in}}$ and $\mathcal{L}_{(0,\vec{x})} \cap M_2^{\text{in}}$ agree with the Lagrangians orthogonal to L_1 and L_2 at $(\vec{x}, 0)$ and $(0, \vec{x})$.

(5.6)

If $|\vec{x}| \leq \delta$ then $\mathcal{L}_{(\vec{x},0)}$ and $(\vec{x}, 0) + L_\Delta$ agree in the region where $|\vec{x} + \vec{y}| \geq \delta$.

(5.7)

If $\delta \leq |\vec{x}| \leq 2$ then $\mathcal{L}_{(\vec{x},0)}$ and $(\vec{x}, 0) + L_\Delta$ agree away from a 4ϵ neighborhood of $L_1 \cup L_2$, and are C^1 -close (with C^1 -distance independent of δ) in this region.

The reader should consult the picture on the right of Figure 7, which illustrates the case $n = 1$. To see the fact that the leaves are parametrized by L , note that, in the second and fourth quadrant, there are two leaves passing through the corner point; these form the non-Hausdorff locus of the space of leaves \mathcal{L} .

In order to construct \mathcal{L} , we choose any monotone function $\kappa: [2, 4] \rightarrow [0, 1]$ which is identically 1 near 2 and identically 0 near 4 and whose derivative is uniformly bounded by 1, and consider the affine Lagrangian planes

$$(5.8) \quad (\vec{x}^0, 0) + \{(\vec{x}, \vec{y}) \mid \kappa(|\vec{x}^0|)x_i = y_i\}.$$

The proof that these planes do not intersect each other in $D^*L_1 \# D^*L_2$, and hence form a foliation, is an elementary computation left to the

reader (see the left side of Figure 7). Note that this foliation satisfies all but Condition (5.5) above. In order to ensure that the leaf corresponding to a vector \vec{x}^0 whose norm lies between 4ϵ and 4 intersects M^{in} in Lagrangian planes orthogonal to L_1 and L_2 , we deform this foliation by affine Lagrangians using an appropriate Hamiltonian flow. The existence of such a flow can be proved, for example, by noting that the region

$$(5.9) \quad 3\epsilon < |\vec{x}| < 4 \text{ and } |\vec{y}| < \epsilon$$

carries a unique function H_0 which vanishes along L_1 and whose restriction to any leaf (5.8) has differential that agrees with the restriction of $-\sum(x_i - x_i^0)dy_i$; we can write its restriction to a leaf explicitly:

$$H_0|_{\mathcal{L}(\vec{x}^0,0)} = \frac{\kappa(|\vec{x}^0|)}{2}|\vec{y}|^2.$$

We define H_T implicitly via the property that it also vanishes along L_1 and that the restriction of H_T to the image of a leaf under the flow ϕ^T generated by the Hamiltonians $\{H_t\}_{t \leq T}$ also has differential agreeing with the restriction of $-\sum(x_i - x_i^0)dy_i$. The reader may easily check that the image of each leaf under ϕ^T is a linear Lagrangian, passing through the same point of L_1 , and that the norm of the restriction of $\sum(x_i - x_i^0)dy_i$ to each leaf decreases with T . Whenever $T = 1$, we find that the image of each leaf is in fact orthogonal to L_1 . We extend H_T to a function on M so that it satisfies the analogous conditions in the region

$$(5.10) \quad 3\epsilon < |\vec{y}| < 4 \text{ and } |\vec{x}| < \epsilon$$

and define the new foliation \mathcal{L} to be the image of (5.8) under ϕ^1 . More precisely, we require H_T to have support contained in a 4ϵ neighborhood of the union of the regions (5.9) and (5.10), and have C^0 and C^1 norms bounded by (say) 2ϵ . These support conditions imply that \mathcal{L} inherits the remaining conditions from the foliation (5.8). Choosing $\epsilon \ll \delta$ implies that the estimate of Lemma 5.1 applies to \mathcal{L} as well:

Lemma 5.2. *The integral of ω over any disc with counterclockwise boundary conditions $(L_2, L_1, \mathcal{L}_{(\vec{x},0)})$ is strictly negative if $\delta \leq |\vec{x}|$. q.e.d.*

5.2. Moduli spaces of caps. In this section, we will discuss moduli spaces of maps from caps to the plumbing of two cotangent bundles. Strictly speaking, the setup we will be using is not exactly the one we will later need to establish the desired results about maps of mushrooms. However, since all the main ideas are already present and the notation is much simpler, we will start here.

Condition (5.2) identifies $Q_{1,2}$ and $Q_{2,1}$ with balls of radius 2 in flat Euclidean space. We assume that the metrics on Q_1 and Q_2 restrict to

this flat metric in U_1 and U_2 . The first step is to impose a final condition on the perturbation $H^{i,j}$ and the Morse functions $f_{i,j}$:

(5.11)

The image of Q_i under the time-1 Hamiltonian flow of $H^{i,j}$ agrees with the graph of the differential of $f_{i,j}$, in a neighborhood of $Q_i \cap Q_j$ which is identified with a (small) disc cotangent bundle of $Q_{i,j}$ using the foliation \mathcal{L}_q . Moreover, the images of Q_1 and Q_2 under the time- t Hamiltonian flow of $H^{i,j}$ are transverse to all leaves \mathcal{L}_q if $|t| \leq 1$. Finally, all critical points of $f_{i,j}$ have distinct values.

Note that when $i = j$, these conditions may easily be achieved by taking any (sufficiently small) generic function $H^{i,i}$. The image of Q^i under such a map may be written uniquely up to a global constant as the graph of the differential of a function $f_{i,i}$. The same strategy works for $i \neq j$ once we observe that L_2 (in \mathbb{C}^n) is the graph of the differential of $|\vec{x}|^2$ as a function on L_1 with respect to the linear foliation $(0, \vec{x}) + L_\Delta$, of which \mathcal{L}_q is a small perturbation. In particular, if we choose $H^{2,1} = H^{1,2} = 0$, we conclude:

Lemma 5.3. *The distance to b grows along gradient flow lines of $f_{1,2}$ (and decreases along gradient lines of $f_{2,1}$) in the subset of U_1 (respectively U_2) where this distance is greater than δ . q.e.d.*

Assumption (5.11) establishes a correspondence between critical points of $f_{i,j}$ and intersection points of the corresponding Lagrangians, which implies:

Corollary 5.4. *There is a bijective correspondence between the generators of $CM^*(f_{i,j})$ and those of $CF^*(Q_i, Q_j)$. q.e.d.*

Floer proved in [6] that, for an appropriate choice of almost complex structure, this correspondence becomes a chain isomorphism. Instead of using this result directly, we shall exhibit a chain equivalence, which will be the first term of an A_∞ functor:

Fix a cap $P \in \mathcal{P}_{\vec{l}}$ for a sequence \vec{l} of length $d + 1$. To each point (z, q) of $\partial P \times Q$, we can associate a Lagrangian $L \subset M$ as follows:

- If s lies on the outgoing segment, we let $L = \mathcal{L}_q$.
- Otherwise, orienting the disc counterclockwise starting at the outgoing segment, $L = Q_j$ if s lies on the j th segment.

Let us assume, as in [17, Section 9g], that we have chosen a universal consistent choice of strip-like ends on the moduli space of discs, i.e., a choice of strip-like end on every stable holomorphic disc which is compatible with gluings for small parameters. Note that the isomorphism $\mathcal{P}_d \cong \mathcal{R}_{d+2}$ therefore determines a choice of strip-like ends on the cap P . In addition, for points sufficiently close to the boundary of $\overline{\mathcal{P}}_d$, we obtain a decomposition into a thin part (obtained as the image of a strip-like

end) and a thick part (the complement). The thin part decomposes as a union of strip-like ends with the images of maps

$$(5.12) \quad [0, 1] \times [-R, R] \rightarrow P.$$

Whenever both boundary components of a strip in the thin part are labeled by a Lagrangian Q_i , the Floer datum determines a perturbation of the $\bar{\partial}$ equation. We extend this choice (which was fixed in the construction of $\text{Fuk}(Q_1, Q_2)$) to the rest of the thin part as follows:

(5.13)

The Floer datum for Q_i and \mathcal{L}_q is given by a fixed almost complex structure J_q and vanishing Hamiltonian.

This is a reasonable assumption because \mathcal{L}_q intersects Q_i transversely at a single point. No breaking of holomorphic strips can therefore occur, so one need perturb neither the complex structure nor the Lagrangians in order to avoid it.

Definition 5.5. A Q -parametrized perturbation datum on P is a pair $(K^{\vec{I}}, J^{\vec{I}})$ where

$$\begin{aligned} K^{\vec{I}} &\in C^\infty(P \times Q, T^*P \otimes \mathcal{H}), \\ J^{\vec{I}} &\in C^\infty(P \times Q, \mathcal{J}). \end{aligned}$$

We require that (i) the restriction of the perturbation data $(K^{\vec{I}}, J^{\vec{I}})$ to the thin part agrees with the labeling Floer data and (ii) if $(z, q) \in \partial P \times Q$ is labeled by the Lagrangian L , then for each tangent vector $\xi \in T_z \partial P$, we have $K^{\vec{I}}(z, q)(\xi)|_L \equiv 0$.

In other words, we choose a family of almost complex structures $J^{\vec{I}}$ on M parametrized by Q and by the points of the disc P , which are required to agree with the choices made in defining the Fukaya category near the incoming marked points, and with J_q near the outgoing points. We also choose a family of 1-forms $K^{\vec{I}}$ on P , parametrized by Q , and with values in the space of functions on M . The 1-forms are also compatible with the Hamiltonian perturbations used to define the Fukaya category.

Note that $K^{\vec{I}}(z, q)(\xi)$ induces a Hamiltonian vector field $Y^{\vec{I}}(z, q)(\xi)$. Given such a datum, for each $q \in Q$ we can therefore consider the perturbed Cauchy-Riemann equation on P :

$$(5.14) \quad (du - Y^{\vec{I}})^{0,1} = 0,$$

where the $(0, 1)$ part is taken with respect to the almost complex structure $J^{\vec{I}}(z, q)$ and the Lagrangian boundary conditions are prescribed by the labels.

Definition 5.6. A Q -parametrized perturbation datum on a singular d -cap P is the choice of a Q -parametrized perturbation datum on each irreducible component.

As the Fukaya category requires counting curves parametrized over the moduli space \mathcal{R}_d , we consider the analogous generalization for caps:

Definition 5.7. *A universal Q -parametrized perturbation datum for caps $(\mathbf{K}^{\mathcal{P}}, \mathbf{J}^{\mathcal{P}})$ is a choice of a Q -parametrized perturbation datum for each singular cap P , smoothly varying with respect to the modulus, and which is compatible with the perturbation data $(\mathbf{K}^{\mathcal{R}}, \mathbf{J}^{\mathcal{R}})$ and the maps (4.8) and (4.9).*

REMARK 5.8. The compatibility condition can be stated explicitly as follows: Every component of a singular cap in the image of the maps (4.8) or (4.9) is a priori equipped with two perturbation data, one coming from being the component of a disc or cap in a factor of the source, the other from being the component of a cap in $\overline{\mathcal{P}}_F$. We require these data to agree.

The proof of the following result follows immediately from the inductive method of proof for the analogous Lemma 9.5 of [17].

Lemma 5.9. *Every perturbation datum on a fixed (smooth) cap may be extended to a universal perturbation datum. q.e.d.*

Let $\vec{p} = \{p_{i_k}\}_{i=1}^d$ denote a sequence of time-1 Hamiltonian chords of H^{i_{k-1}, i_k} points starting at Q_{i_k} and ending on $Q_{i_{k-1}}$. In addition, we will denote the point of intersection $\mathcal{L}_q \cap Q_i$ by q_i .

Definition 5.10. *The moduli space of caps with inputs \vec{p} is the space $\mathcal{P}(\vec{p}) = \coprod_{q \in Q} \mathcal{P}_q(\vec{p})$ of solutions to the perturbed Cauchy Riemann equation (5.14) with incoming marked points mapping to \vec{p} and outgoing marked points mapping to q_{i_0} and q_{i_d} which project to the same point $q \in Q$.*

In Section 7.2 we prove transversality for the moduli space of caps:

Lemma 5.11. *For a generic choice of universal consistent Q -parametrized cap perturbation data, $\mathcal{P}(\vec{p})$ is a smooth manifold. More generally, any stratum of $\overline{\mathcal{P}}(\vec{p})$ consisting of a single cap (and an arbitrary number of discs) has the expected dimension, and, if $N \subset Q$ is a submanifold, the projection map $\mathcal{P}(\vec{p}) \rightarrow Q$ (at every smooth point of the source) is generically transverse to N on every such stratum.*

We shall also prove compactness in Section 6, by appropriately generalizing Lemma 5.1:

Lemma 5.12. *Whenever $i_0 = i_d$ or $(i_0, i_d) = (2, 1)$, the Gromov bordification $\overline{\mathcal{P}}(\vec{p})$ is compact. In the second case, $\mathcal{P}_q(\vec{p})$ is in addition empty whenever q lies more than 4ϵ away from b .*

If $(i_0, i_d) = (1, 2)$, then the Gromov bordification of the moduli space of caps with boundaries on leaves corresponding to points within 4ϵ of q is compact.




Figure 8. The deepest stratum of \mathcal{C}_d .

5.3. Mushroom maps. Given the seeming complexity of the definitions coming up, a word of explanation is in order. Naively, a mushroom map is simply a holomorphic map defined on each cap together with a gradient tree satisfying certain compatibility conditions. Let us fix the homeomorphism type of a tree S with r inputs, and a sequence of integers d_j such that $\sum d_j = d$. Assume we have fixed in addition, a sequence of labels \vec{I} of length d , a sequence $\vec{p} = \vec{p}^1 \cup \dots \cup \vec{p}^r$ of Hamiltonian chords with endpoints on the Lagrangians $Q_{i_{k-1}}$ and Q_{i_k} , and a critical point x_0 of f_{i_0, i_d} . Each sequence \vec{p}^k determines a subsequence \vec{I}^k of \vec{I} , whose elements we denote $\{i_j^k\}_{j=1}^{d_k}$.

We would have liked to define the space of mushroom maps with inputs \vec{p} , output x_0 , and fixed topological type as a fiber product

$$(5.15) \quad \mathcal{S}_S(x_0, Q_{i_0=i_0^1, i_0^2}, \dots, Q_{i_r=i_r^1, i_r^2, \dots, i_r^r}) \times_{Q^r} (\mathcal{P}(\vec{p}^1) \times \dots \times \mathcal{P}(\vec{p}^r))$$

where $\vec{I}^k = \{i_j^k\}$ is the sequences of labels for the incoming segments of an element of $\mathcal{P}(\vec{p}^j)$. Whenever S has no incoming leaves of vanishing length, one may easily choose perturbation data on the moduli space of shrubs and caps once and for all, so that all such fiber products are transverse (we shall see later how the lack of Hausdorffness is circumvented). However, there is no way to achieve transversality in this way for the stratum shown in Figure 8 whenever some of the caps have the same inputs.

The main observation (we used a similar trick in achieving regularity for the moduli space of cascades in [3]) is that there is no reason for the perturbation data on the various caps of a fixed mushroom to be defined in exactly the same way. In particular, even if all the caps of the mushroom shown in Figure 8 have marked points labeled by the same Hamiltonian chords, our setup will allow up to choose different perturbation data for each of them, thereby making transversality easy to achieve. The rest of this section implements this idea, which we shall return to repeatedly (see in particular Remark 5.20 and the discussion following Lemma 7.14)

We first consider a simplest case of a mushroom with one input whose incoming boundary segments are labeled (i_0, i_1) .

Definition 5.13. A Q -parametrized perturbation datum on the mushroom with one input (which is unique) is the choice of (i) a family of almost complex structures in \mathcal{J} parametrized by a disc with three punctures, agreeing with the respective complex structure at the end and (ii) a closed 1-form γ on the cap with three punctures whose pull-back under the positive end with boundary labels (Q_{i_0}, Q_{i_1}) and the negative end with boundary (Q_{i_0}, \mathcal{L}_q) agrees with dt , and which vanishes near the remaining end.

In this case, we fix the Hamiltonian perturbation datum as

$$(5.16) \quad K^{i_0, i_1} \equiv H^{i_0, i_1} \otimes \gamma.$$

If we have more inputs, we shall also require a perturbation of the stem. Recall that a labeling \vec{I} of the inputs of a mushroom C is equivalent to the data of compatible labelings \vec{I}^k of all caps, and \vec{R} of the stem. In particular, each edge of the stem is equipped with a gradient flow equation which we shall have to perturb. Unlike the usual setup, we shall allow the Floer data on the ends meeting the outgoing segment of a cap to vary depending on the modulus. This does not cause any difficulties, as no breaking of holomorphic strips can happen at these ends.

Definition 5.14. A Q -parametrized perturbation datum on a singular mushroom C is the choice of (i) a gradient perturbation datum on the stem of C and (ii) a Q -parametrized perturbation datum on each (possibly singular) cap such that we have the following:

(5.17) *The Floer datum for $(\mathcal{L}_q, Q_{i,d})$ and (Q_{i_0}, \mathcal{L}_q) is given by the constant choice of Condition (5.13) with the following exception: If the j th cap has exactly 2 incoming segments with Lagrangian labels (Q_{i_0}, Q_{i_1}) , then the Hamiltonian perturbation associated to the end (Q_{i_0}, \mathcal{L}_q) is given by multiplying H^{i_0, i_d} by a non-negative constant which is bounded by 1. This constant vanishes whenever the j th incoming leaf of the stem is short (say, has length less than 10) and equal 1 whenever it is sufficiently long.*

REMARK 5.15. The perturbation of the gradient flow need not be compatible with the simplicial subdivision \mathcal{Q} as it was in Section 3. However, we still require that the gradient flow be transverse to the boundary of $Q_{i,j}$ (and pointing in the appropriate direction).

REMARK 5.16. The constraint (5.17) is designed to achieve two things: (i) whenever a mushroom with a single input breaks off, we need the perturbation datum to agree with the one chosen in Definition 5.13 and (ii) whenever two successive incoming edges of the stem have vanishing length, we require vanishing Floer data in order for the gluing theorem to hold.

Of course, the perturbation datum is not strictly speaking parametrized by Q anymore, but rather by Q^r . For the next definition, we fix, as before a sequence $\vec{p} = (p_1, \dots, p_d)$ of Hamiltonian chords with endpoints on Q_1 and Q_2 and a critical point x_0 of f_{i_0, i_d} .

Definition 5.17. *A singular mushroom map Ψ with inputs \vec{p} and output x_0 is a collection of maps*

$$\begin{aligned}\Psi_S &: S(C) \rightarrow Q, \\ \Psi_j &: P_j(C) \rightarrow M,\end{aligned}$$

where Ψ_S is a (singular) perturbed gradient tree (as in Definition 3.6) and Ψ_j are finite energy solutions to the perturbed Cauchy-Riemann equation (5.14) on the relevant (possibly singular) caps satisfying the following conditions:

- 1) *If the k th incoming vertex lies on $P_j(C)$, then Ψ_j converges on the corresponding end to p_k , and the incoming boundary segments of $P_j(C)$ are mapped to the labeling Lagrangians.*
- 2) *The outgoing arc on $P_j(C)$ is mapped to $\mathcal{L}_{\Psi_S(v_j)}$, where v_j is the j th incoming leaf of the stem.*

Note that the last condition and the removal of singularities theorem forces the outgoing marked points of the j th cap to be mapped to the lifts of $\Psi_S(v_j)$ to the appropriate manifold Q_1 or Q_2 .

The next result is the main tool in establishing compactness for maps of mushrooms:

Lemma 5.18. *The image of every cap of a singular mushroom map is contained in a compact subset of M , and every edge of a stem labeled by (i, j) with $i \neq j$ is mapped to the ball of radius δ in $Q_{i,j}$. In particular, if the distance between q and b is greater than δ , there is no mushroom map admitting a cap whose first and last labels differ and whose boundary condition on the outgoing arc is \mathcal{L}_q .*

Proof. The proof is a combination of Lemma 5.12 and the proof of Lemma 2.19. We focus on the case of a cap $P_j(C)$ whose first segment (in the counterclockwise order as usual) is mapped to Q_1 , and penultimate one to Q_2 (Lemma 5.12 excludes the other possibility).

As in the proof of Lemma 2.19, we find that either (i) there is a descending arc on the stem $S(C)$, starting at the j th incoming leaf of $S(C)$ and ending on the outgoing edge, all of whose edges are labeled $(1, 2)$ or (ii) there is an arc from the j th incoming leaf of $S(C)$ to another leaf meeting an incoming cap $P_{j'}(C)$ which can be broken into a descending arc with edges labeled $(1, 2)$, followed by an ascending arc with edges labeled by $(2, 1)$. Since the perturbed negative gradient flow along an edge labeled by $(1, 2)$ points outward (and inward for edges labeled by $(2, 1)$), we conclude that the critical point of $f_{1,2}$ which is the output of

our mushroom map (or the image of the cap $P_{j'}(C)$ under the projection to $Q_{2,1}$) lies outside the ball of radius δ about the base point in $Q_{1,2}$ (or in $Q_{2,1}$ in the second case). This contradicts our assumption (5.11) on the gradient flow $f_{1,2}$, or Lemma 5.12 which we can apply to the cap $P_{j'}(C)$. q.e.d.

Following the construction for discs and trees, we shall consider perturbation data varying with the modulus of the mushroom:

Definition 5.19. *A universal consistent Q -parametrized perturbation datum for mushrooms is a choice $(\mathbf{X}^C, \mathbf{K}^C, \mathbf{J}^C)$ of Q -parametrized perturbation data for each singular mushroom C , varying smoothly in each stratum of $\bar{\mathcal{C}}$, and which is compatible with (i) the maps (4.7) and the perturbation data $(\mathbf{K}^R, \mathbf{J}^R)$ used to define $\text{Fuk}(Q_1, Q_2)$ and (ii) the maps (4.6) and the perturbation datum \mathbf{X}^T used to define $\text{Morse}(Q_1, Q_2)$. Moreover, whenever the incoming leaf of the stem meeting a given cap P_j has length smaller than 1, we require that the Floer datum on the ends adjacent to its outgoing segment to be the constant ones fixed in (5.13).*

REMARK 5.20. Had we naively pulled back the perturbation datum for a mushroom using the projections of components of \mathcal{C}_d to a product as in Equation (4.5), we would not have been able to achieve transversality in the situation pictured in Figure 8. However, while the stratum labeled by the topological type of such mushrooms lies on the boundary of the Gromov compactification of a component of \mathcal{C}_d , it in fact lies in the interior of $\bar{\mathcal{C}}_d$. In particular, the choice of perturbation data for a fixed d on this part of the moduli space is essentially independent from the choices made for smaller value of d . This implies that the choice of datum of each cap is unconstrained by the other choices, thereby allowing us to achieve transversality as explained in the discussion following Lemma 7.14.

Given such a perturbation datum, we can define the moduli spaces which allows us to construct a functor from $\text{Morse}(Q_1, Q_2)$ to $\text{Fuk}(Q_1, Q_2)$.

Definition 5.21. *The moduli space of mushroom maps*

$$\mathcal{C}(x_0; \vec{p})$$

is the space of maps Ψ from a non-singular mushroom C with incoming marked points \vec{p} and outgoing marked point x_0 . The space of stable maps from a possibly singular mushroom is denoted

$$\bar{\mathcal{C}}(x_0; \vec{p}).$$

The topology on $\bar{\mathcal{C}}(x_0; \vec{p})$ is defined as usual using the Gromov C^∞ topology for the underlying caps and stem. Moreover, this space carries a natural stratification coming from the stratifications on $\bar{\mathcal{P}}$ and $\bar{\mathcal{S}}$

(and possible breakings of gradient flow lines between critical points, or holomorphic strips at the incoming ends). We state the main result of Section 7.2.2:

Lemma 5.22. *For a generic choice of perturbation data, all strata of the moduli space $\overline{\mathcal{C}}(x_0; \vec{p})$ are of the expected dimension.*

The proof of this next result is given in Section 7.2.3 and requires a gluing result in codimension 1 for moduli spaces of caps:

Proposition 5.23. *Given generic perturbation data, $\overline{\mathcal{C}}(x_0; \vec{p})$ is a compact 1-dimensional manifold with boundary whenever*

$$(5.18) \quad \deg(x_0) = 1 - d + \sum \deg(\vec{p}).$$

5.4. Construction of the functor. We begin by defining the linear term of the functor \mathcal{G} promised in the beginning of this section. To do this, we first observe that the index theorem (say, Equation (12.2) in [17]) gives a natural isomorphism

$$(5.19) \quad \lambda(\mathcal{P}_q(p_1)) \cong \mathfrak{o}_{q_{i_0}}^\vee \otimes \mathfrak{o}_{p_1}^\vee \otimes \mathfrak{o}_{q_{i_1}} \cong \mathfrak{o}_{p_1}^\vee$$

where $\mathcal{P}_q(p_1)$ is the moduli space of caps projecting to a fixed point q in Q , and we are writing q_i for the intersection point between Q_i and \mathcal{L}_q . It follows immediately that the determinant bundle of the moduli space of caps admits a natural isomorphism

$$(5.20) \quad \lambda(\mathcal{P}(p_1)) \cong \mathfrak{o}_{p_1}^\vee \otimes \lambda(Q).$$

Given a mushroom Ψ in $\mathcal{C}(x_0, p_1)$, the description of the moduli space of such mushrooms as the fiber product of $\mathcal{P}(p_1)$ with $W^s(x_0)$ over Q gives an isomorphism

$$(5.21) \quad \lambda(\mathcal{C}(x_0, p_1)) \otimes \lambda(Q) \cong \lambda(W^s(x_0)) \otimes \lambda(\mathcal{P}(p_1)).$$

Specializing to the case $\deg(x_0) = \deg(p_1)$, and using the previous result for $\lambda(\mathcal{P}(p_1))$ and the definition of \mathfrak{o}_{x_0} , we conclude that every mushroom Ψ determines an isomorphism

$$(5.22) \quad \mathfrak{o}_{p_1} \cong \mathfrak{o}_{x_0},$$

and hence a map

$$(5.23) \quad \mathcal{G}^\Psi: |\mathfrak{o}_{p_1}| \rightarrow |\mathfrak{o}_{x_0}|$$

which is the contribution of Ψ to the linear part of the functor \mathcal{G} :

$$\begin{aligned} \mathcal{G}^1: CF^*(Q_i, Q_j) &\rightarrow CM^*(f_{i,j}) \\ [p_1] &\mapsto (-1)^{n \deg(p_1)} \sum_{\substack{\deg(x_0) = \deg(p_1) \\ \Psi \in \mathcal{C}(x_0, p_1)}} \mathcal{G}^\Psi([p_1]). \end{aligned}$$

Ignoring signs, the fact that \mathcal{G}^1 is a chain map is a consequence of Proposition 5.23. Indeed, considering p_1 and x_0 such that

$$(5.24) \quad \deg(x_0) = \deg(p_1) + 1,$$

the boundary strata of $\overline{\mathcal{C}}(x_0, p_1)$ consist of either (i) a critical point x_1 of $f_{i,j}$ and an element of $\mathcal{T}(x_0, x_1) \times \mathcal{C}(x_1, p_1)$ (i.e., a rigid mushroom with input p_1 and output x_1 together with a rigid gradient trajectory from x_1 to x_0) or (ii) an element of $\mathcal{C}(x_0, p_0) \times \mathcal{R}(p_0, p_1)$ for a time-1 chord p_0 from Q_i to Q_j for the Hamiltonian $H^{i,j}$. These two cases correspond to the two sides of the equation

$$(5.25) \quad \mathcal{G}^1 \circ \mu_1^{\text{F}} = \mu_1^{\text{S}} \circ \mathcal{G}^1.$$

To prove the correctness of the signs, we must compare the natural isomorphism (5.21), with a product isomorphism coming from its boundary strata. The case (i) above is virtually indistinguishable from the analogous situation for the moduli spaces of shrubs interpolating between simplicial and Morse cochains, analysed in the discussion surrounding Equation (3.28). The result of that discussion is that there is a sign difference given by the parity of $\deg(p_1)$ between the product orientation, and the one induced from the interior.

In the other case, we first note that a rigid strip gives a canonical isomorphism

$$(5.26) \quad \lambda(\mathbb{R}) \cong \mathfrak{o}_{p_0} \otimes \mathfrak{o}_{p_1}^{\vee}$$

with the left-hand side corresponding to translation. Taking the tensor product with Equation (5.21) with p_1 replaced by p_0 , and $TC(x_0, p_0)$ trivialized canonically by virtue of being rigid, we obtain an isomorphism

$$(5.27) \quad \lambda(Q) \otimes \lambda(\mathbb{R}) \cong \lambda(W^s(x_0)) \otimes \mathfrak{o}_{p_0}^{\vee} \otimes \lambda(Q) \otimes \mathfrak{o}_{p_0} \otimes \mathfrak{o}_{p_1}^{\vee}.$$

Since gluing the strip to the cap at p_0 gives an identification $\mathbb{R} \cong T\overline{\mathcal{C}}(x_0, p_1)$ with the positive direction in \mathbb{R} pointing inward, we recover the isomorphism of (5.21), multiplied by -1 .

To conclude that (5.25) holds, we must keep in mind that there are signs of $(-1)^{\deg(p_1)}$ and $(-1)^n$ in the definitions of the differential in the Floer and Morse theories, as well as a sign in our formula for \mathcal{G}^1 , which contribute $n \deg(p_1)$ in one case, and $n \deg(p_1) + n$ in the other.

Proposition 5.24. *The map \mathcal{G}^1 is a chain isomorphism.*

Proof. We shall use the energy filtration to see that, if we order the critical points by energy, the matrix for \mathcal{G}^1 is upper triangular with ones along the diagonal, and hence is invertible.

We first show that $\mathcal{C}(x_0; p_1)$ is empty whenever

$$(5.28) \quad f_{i_0, i_1}(x_1) < f_{i_0, i_1}(x_0)$$

where x_1 is the critical point corresponding to p_1 in Corollary 5.4. Assume by contradiction that a mushroom map Ψ with these asymptotic conditions exists. Let q denote the point of Q such that the cap of Ψ has a boundary segment mapping to \mathcal{L}_q , and write q_{i_0} for the lift of this point to Q_{i_0} . Since the stem is a descending flow line of f_{i_0, i_1} , we conclude that

$$(5.29) \quad f_{i_0, i_1}(x_0) \leq f_{i_0, i_1}(q),$$

and hence, as we are assuming (5.28), we find that

$$(5.30) \quad f_{i_0, i_1}(x_1) < f_{i_0, i_1}(q).$$

We claim that this inequality contradicts positivity of energy for the cap. Indeed, an application of Stokes's theorem analogous to Lemma 5.1, the closedness of the 1-form in Definition 5.13 implies that

$$(5.31) \quad \int |d\Psi_P|^2 = \mathcal{A}_0(q_{i_1}) + \mathcal{A}_{H^{i_0, i_1}}(q_{i_0}) - \mathcal{A}_{H^{i_0, i_1}}(p)$$

where q_{i_0} and q_{i_1} now stand for the intersection points of \mathcal{L}_q with the respective Lagrangian and \mathcal{A} for the appropriate action functional for the Hamiltonian perturbations fixed in Definition 5.13. Using the primitive θ_{i_0} which vanishes on Q_{i_0} and \mathcal{L}_q , we compute that $\mathcal{A}_0(q_{i_1})$ vanishes, while $\mathcal{A}_{H^{i_0, i_1}}(q_{i_0})$ and $\mathcal{A}_{H^{i_0, i_1}}(p)$ are, respectively, equal to $-f_{i_0, i_1}(q)$ and $-f_{i_0, i_1}(x_1)$. As the left-hand side of Equation (5.31) is non-negative, we conclude, as desired, that $\mathcal{C}(x_0; p_1)$ is empty whenever the inequality (5.28) holds.

Applying this energy analysis to the case p is the Hamiltonian chord from Q_{i_0} to Q_{i_1} corresponding to the critical point x of f_{i_0, i_1} , we find that the only solutions in this case must have constant cap and stem (any non-constant solution must carry some energy). The proof that there are ones along the diagonal and hence that \mathcal{G}^1 is a chain isomorphism now follows from the fact that such constant solutions are regular. q.e.d.

The higher-order analogues of this map are defined in essentially the same manner as the maps \mathcal{F}^d which define a functor from Floer to Morse theory. Whenever $\deg(x_0) = 1 - d + \sum_i \deg p_i$, we may assign to each tree $\Psi \in \mathcal{C}(x_0, \vec{p})$ an isomorphism

$$\mathfrak{o}_{p_1} \otimes \cdots \otimes \mathfrak{o}_{p_d} \rightarrow \mathfrak{o}_{x_0}$$

and a sign $\ddagger(\Psi)$. We discuss a set of possible choices for the isomorphism and the sign in Section 8.4. We define the d th higher term of \mathcal{G} ,

$$\mathcal{G}^d: CF^*(Q_{i_{d-1}}, Q_{i_d}) \otimes \cdots \otimes CF^*(Q_{i_0}, Q_{i_1}) \rightarrow CM^*(f_{i_0, i_d}),$$

to be a sum of the maps \mathcal{G}^Ψ induced by Ψ on orientation lines, multiplied by the appropriate sign:

$$[p_d] \otimes \dots \otimes [p_1] \mapsto \sum_{\substack{\Psi \in \mathcal{C}(x_0, \bar{p}) \\ \deg(x_0) = \sum_k \deg p_k + d - 1}} (-1)^{\ddagger(\Psi)} \mathcal{G}^\Psi([p_f] \otimes \dots \otimes [p_1]).$$

Modulo signs, the analysis of the boundary strata of 1-dimensional moduli spaces of mushrooms given in Section 7.2.3 (i.e., the proof of Proposition 5.23) immediately implies:

Proposition 5.25. *The maps \mathcal{F}^d satisfy the A_∞ equation for a functor:*

$$(5.32) \quad \sum_{\substack{r \\ d_1 + \dots + d_r = d}} \mu_r^M(\mathcal{F}^{d_1} \otimes \dots \otimes \mathcal{F}^{d_r}) = \sum_{\substack{d_1 + d_2 = d + 1 \\ i < d_1}} (-1)^{\mathfrak{K}_i} \mathcal{F}^{d_1}(\text{id}^{d_1 - i - 1} \otimes \mu_{d_2}^F \otimes \text{id}^i).$$

q.e.d.

6. Compactness for caps

The goal of this section is to prove Lemma 5.12. We start in the first section by discussing a general result which constrains holomorphic curves to remain within compact domains, which we shall then apply in the second section to prove the desired statement:

6.1. Generalities. We begin with a generalization of Lemma 7.2 of [3], which precludes holomorphic curves from escaping from a domain. Let us assume therefore that we are given an almost complex structure I_W on a manifold W , and a 1-form α whose differential is non-negative on every complex plane. Let W^{in} be a (codimension 0) submanifold with boundary such that $\ker(\alpha|_{\partial W^{\text{in}}})$ is a complex distribution (in particular, α does not identically vanish near ∂W^{in}), and

$$(6.1) \quad \begin{aligned} &\text{if } \alpha(X) = 1 \text{ for } X \in T\partial W^{\text{in}} \text{ then (i) } I_W X \text{ is inward pointing and} \\ &\text{(ii) } d\alpha \text{ is positive on the complex plane spanned by } X. \end{aligned}$$

Assuming moreover that we are given a submanifold $R \subset W - W^{\text{in}}$ transverse to ∂W^{in} such that

$$(6.2) \quad \begin{aligned} &\alpha \text{ restricts to an exact form on } R \text{ whose primitive is constant on} \\ &R \cap \partial W^{\text{in}}, \end{aligned}$$

we conclude:

Lemma 6.1. *Every I_W holomorphic curve in $W - \text{int}(W^{\text{in}})$ with boundary on $R \cup \partial W^{\text{in}}$ which intersects ∂W^{in} is contained therein.*

Proof. Since the proof is essentially indistinguishable from that of Lemma 7.2 in [3] we shall only sketch it: The assumption that $d\alpha$ is non-negative on holomorphic planes implies that its integral over a holomorphic curve is non-negative. Moreover, given that $d\alpha$ does not vanish on a plane transverse to ∂W^{in} this integral must be strictly positive for every holomorphic curve which intersects ∂W^{in} but is not contained in it. In particular, Stokes's theorem implies that the integral of α over the boundary of such a holomorphic curve must be positive. However, the contribution of the subset mapping to R vanishes by Condition (6.2), while the contribution of the subset mapping to ∂W^{in} is non-positive by Condition (6.1). q.e.d.

Still requiring that $\alpha|_R$ be exact, we now assume that the primitive is only locally constant on $R \cap \partial W^{\text{in}}$. In this case, let us consider maps

$$u: \Sigma \rightarrow W - \text{int}(W^{\text{in}})$$

such that whenever $\gamma: [0, 1] \rightarrow \partial\Sigma$ is a parametrization of a boundary segment mapping to R which is compatible with the complex orientation,

(6.3)

the difference between the values of a primitive of $\alpha|_R$ at $u(1)$ and $u(0)$ is strictly positive.

The previously sketched argument in fact proves:

Lemma 6.2. *There is no holomorphic curve in $W - \text{int}(W^{\text{in}})$ with boundary on $R \cup \partial W^{\text{in}}$ for which Condition (6.3) holds.* q.e.d.

6.2. Proof of compactness. If in the discussion of the previous section, α is a Liouville form whose restriction to ∂W^{in} is a contact form, then Condition (6.1) essentially says that I_W is of contact type near ∂W^{in} , and Condition (6.2) that R is exact with Legendrian boundary (and in addition admits a primitive for the Liouville form which strictly vanishes on this boundary). In this case, we exactly recover Lemma 7.2 of [3], and we may conclude half of Lemma 5.12:

Lemma 6.3. *If the distance between q and b is greater than δ then $\overline{\mathcal{P}}_q(\vec{x})$ is compact whenever $i_0 = i_d$ and empty whenever $(i_0, i_d) = (2, 1)$.*

Proof. As usual, it suffices to prove the existence of a compact subset of M (in this case M^{in}) containing the image of every element of $\mathcal{P}_q(\vec{x})$, so that the usual proof of Gromov compactness applies. Let $u: S \rightarrow M$ be such a map, and denote by Σ the inverse image of $M - \text{int}(M^{\text{in}})$. If $i_0 = i_d$, then the hypothesis of Lemma 6.1 holds if α is a primitive for the symplectic form which agrees with the natural primitive near M^{in} , and the image of $u|_\Sigma$ is contained in ∂M^{in} ; hence the image of u is contained in M^{in} . If $(i_0, i_d) = (2, 1)$, then the computation of Lemma 5.2 shows that Condition (6.3) holds, so that Lemma 6.2 implies the desired result. q.e.d.




Figure 9

Note that while the almost complex structure is of contact type near any point in ∂M^{in} , the Lagrangians \mathcal{L}_q may not intersect it in Legendrian submanifolds if q is within δ of b . We shall construct a different submanifold M^{mi} which will play the same role as M^{in} did in the previous Lemma (this manifold was mentioned in Condition (2.43) and we are finally about to define it). Its local model in \mathbb{C}^n is a smoothing of the manifold with corners

$$(6.4) \quad \{(\vec{x}, \vec{y}) \mid |\vec{x}|^2 \leq 1/4 \text{ or } |\vec{y}|^2 \leq 1/4\} \cup \{(\vec{x}, \vec{y}) \mid |\vec{y}|^2 + |\vec{x}|^2 \leq 2 \text{ and } |\vec{y} + \vec{x}|^2 \leq 4 - \delta\}$$

(see Figure 9). To understand the construction, it is helpful to keep in mind that the sphere of radius 2 about the origin meets the hypersurface $|\vec{y} + \vec{x}|^2 = 4$ precisely on its intersection with the diagonal. Hence this intersection takes place far away from the hypersurfaces $|\vec{x}|^2 = 1/4$ and $|\vec{y}|^2 = 1/4$. In particular, this manifold has only codimension 2 corners, and the smoothing construction which produces M^{mi} is entirely local. We leave the details of the smoothing to the reader, requiring only that

if the distance between b and q is smaller than δ , the intersection of the leaf \mathcal{L}_q with ∂M^{mi} agrees with its intersection with the hypersurface $|\vec{y} + \vec{x}|^2 = 4 - \delta$.

Note that $\{(\vec{x}, \vec{y}) \mid |\vec{y} + \vec{x}|^2 \leq 4 - \delta\}$ can be naturally identified with a disc cotangent bundle of the Lagrangian

$$L_{-\Delta} = (\vec{x}, -\vec{x}) \subset \mathbb{C}^n$$

such that affine planes parallel to L_{Δ} correspond to cotangent fibers. In particular, the leaves \mathcal{L}_q (whenever q is close to b) intersect M^{mi} in Legendrians. With this at hand, we complete this section with the following proof:

Proof of Lemma 2.19. Given that we have already established Lemma 6.3 it remains to show that caps with boundary conditions along \mathcal{L}_q cannot escape ∂M^{mi} whenever q is within δ of b . Condition (2.43) implies that all almost complex structures are of contact type near ∂M^{mi} , and that there are no Hamiltonian perturbations in $M - M^{\text{mi}}$; as in Lemma 6.3, we may directly apply Lemma 6.1 to conclude the inverse image of the complement of M^{mi} is empty. q.e.d.

7. Transversality

7.1. Transversality and gluing for gradient trees. In this section, we explain the proofs of results from Sections 2.2 and 3.

7.1.1. Transversality for Stasheff trees. In order to prove transversality, we must study the moduli space $\overline{\mathcal{T}}_d$ more carefully. First, we recall that its boundary strata are labeled by trees with d inputs

$$(7.1) \quad T^\infty.$$

An element of $\overline{\mathcal{T}}_d$ lies in the cell labeled by T^∞ if the result of collapsing all its finite edges yields a (ribbon) tree combinatorially isomorphic to T^∞ (see [17] which gives a detailed discussion of the parallel case in which the Stasheff polytopes are modeled by moduli spaces of holomorphic discs). The space $\overline{\mathcal{T}}_d$ also has a cell decomposition with cells $\mathcal{T}^{T \rightarrow T^\infty}$ labeled by surjective maps of ribbon trees

$$(7.2) \quad T \rightarrow T^\infty$$

where both T and T^∞ are isomorphism classes of trees with d inputs and such that the inverse image of every vertex is a connected tree. The stratum determined by such a map consists of trees which are homeomorphic to T and whose finite edges are those collapsed in T^∞ . In particular, the set of infinite edges in such a stratum is fixed, so the structure of a metric tree is determined by the lengths of the remaining edges, and we have a canonical isomorphism

$$(7.3) \quad \mathcal{T}^{T \rightarrow T^\infty} \cong (0, +\infty)^{|\mathcal{E}(T)| - |\mathcal{E}(T^\infty)|},$$

with closure (allowing edges to reach length 0 or ∞):

$$(7.4) \quad \overline{\mathcal{T}}^{T \rightarrow T^\infty} \cong [0, +\infty)^{|\mathcal{E}(T)| - |\mathcal{E}(T^\infty)|}.$$

The cell structure on $\overline{\mathcal{T}}^{T \rightarrow T^\infty}$ induced from its inclusion in $\overline{\mathcal{T}}_d$ is compatible with the natural cell structure on $[0, +\infty)^{|\mathcal{E}(T)| - |\mathcal{E}(T^\infty)|}$. Given a sequence \vec{I} of elements of the set $\{1, 2\}$, there is an obvious meaning to the moduli space

$$(7.5) \quad \overline{\mathcal{T}}_{\vec{I}}^{T \rightarrow T^\infty}$$

of labeled Stasheff trees which lie in the appropriate stratum.

We shall now formally describe what we mean by a smooth choice of perturbation data. Let us once and for all fix a smooth structure on $[0, +\infty]$ using the identification

$$(7.6) \quad e^{-x}: [0, +\infty] \cong [0, 1].$$

In particular, we obtain a smooth structure on $\overline{\mathcal{T}}^{T \rightarrow T^\infty}$. Note that every flag in T (a pair (e, v) with v a vertex of the edge e) and the metric g_T determine a unique orientation and length-preserving map

$$(7.7) \quad e \rightarrow \mathbb{R}$$

taking v to the origin and e to the appropriate segment on either side of the origin. In particular, this flag determines a polyhedral subset in

$$(7.8) \quad [0, +\infty]^{|\mathcal{E}(T)| - |\mathcal{E}(T^\infty)|} \times \mathbb{R}$$

whose fiber is the image of the segment e under (7.7).

Definition 7.1. *A choice of perturbation data for each element of $\overline{\mathcal{T}}_{\vec{I}}$ is said to be **smooth** if, given a flag (e, v) of a tree T , the perturbation data for the edge e of metric trees in $\overline{\mathcal{T}}_{\vec{I}}^{T \rightarrow T^\infty}$ extends to a smooth map*

$$(7.9) \quad \tilde{\mathbf{X}}_{(e,v)}^{T \rightarrow T^\infty}: [0, +\infty]^{|\mathcal{E}(T)| - |\mathcal{E}(T^\infty)|} \times \mathbb{R} \rightarrow C^\infty(TM).$$

Recall that a universal perturbation is a smooth choice of perturbation data for all labeled Stasheff trees, which is moreover consistent in the sense that the perturbation data for a singular tree is given by the data of its components.

Lemma 7.2. *Given any sequence \vec{I} , every smooth and consistent choice of perturbation data on the union of*

$$(7.10) \quad \coprod_{|\vec{I}'| < |\vec{I}|} \overline{\mathcal{T}}_{\vec{I}'}$$

with the k -skeleton of $\overline{\mathcal{T}}_{\vec{I}}$ extends smoothly and consistently to the $k+1$ skeleton of $\overline{\mathcal{T}}_{\vec{I}}$. Moreover, the restriction map from the space of all perturbation data to the space of perturbation data for a given singular labeled Stasheff tree (T, g_T) is surjective.

Proof. Since the perturbation data have been fixed on the lower-dimensional moduli spaces it suffices to extend the perturbation datum to interior cells. Fix a tree T labeling such a cell, (in this case T^∞ is a tree with a unique $(d+1)$ -valent vertex, and all interior edges are collapsed, so we may drop this data from the notation). By assumption, perturbation data have been chosen for the boundary of the cell labeled by T . If we fix an interior edge e , writing (e, v_\pm) for the two flags it lies on, then the existence of a smooth extension of a function from the boundary of

$$[0, +\infty]^{|\mathcal{E}(T)| - |\mathcal{E}(T^\infty)|} \times \mathbb{R}$$

is obvious since this is a manifold with corners (we proceed by induction, extending over small neighborhoods of the skeleta). It suffices therefore to make sure that the extension can be chosen so the maps associated to the two flags containing e are both smooth. Writing t_e for the coordinate of the factor of $[0, +\infty]^{|\mathcal{E}(T)|-|\mathcal{E}(T^\infty)|}$ corresponding to the edge e , we find that the only potentially interesting extension problem occurs at the boundary facets $t_e = 0, +\infty$. Since the map $(t_e, y) \mapsto (t_e, y + t_e)$ is smooth on \mathbb{R}^2 , and takes the polyhedron associated to (e, v_-) to the one associated to (e, v_+) , smoothness of the extensions $\tilde{\mathbf{X}}_{(e, v_-)}^{T \rightarrow T^\infty}$ and $\tilde{\mathbf{X}}_{(e, v_+)}^{T \rightarrow T^\infty}$ are equivalent near $t_e = 0$.

It remains to check the existence of an extension from $t_e = +\infty$. The key fact is that the perturbation datum is required to be compactly supported on every edge. This defines a naive way of gluing the perturbation data associated to the flags (e, v_\pm) whenever $t_e \gg 0$ by “pregluing” the two half-infinite intervals as follows: if we identify the edge with the interval $[0, t_e]$, then the perturbation datum on $[t_e/3, 2t_e/3]$ is set to vanish, while those on $[0, t_e/3]$ and $[2t_e/3, t_e]$ are obtained by pulling back the perturbation datum from the natural inclusions in $[0, +\infty)$ and $[-\infty, 0]$ (the second is obtained after translation by $-t_e$). A simple computation shows that the resulting perturbation datum extends smoothly with respect to both coordinate systems, and the proof proceeds by induction using cutoff functions, proving the first part of the result. The second part follows easily by observing that once one has constructed a perturbation datum over the boundary of a cell, it can be arbitrarily modified at any point in its interior. q.e.d.

The next technical result is the infinitesimal version of the fact that perturbing of the gradient flow equation on a bounded subset of an edge integrates to an essentially arbitrary diffeomorphism (see the discussion preceding Definition 2.6).

Lemma 7.3. *For each gradient tree $\psi: (T, g_T) \rightarrow Q$, x a point in the appropriate ascending or descending manifold, and tangent vector X at $\text{ev}_k(x)$, there exists a 1-parameter family of universal perturbation data \mathbf{X}_τ^T whose evaluation maps $\text{ev}_{k'}^\tau$ are independent of τ for $k' \neq k$, and such that*

$$(7.11) \quad \left. \frac{d \text{ev}_k^\tau(x)}{d\tau} \right|_{\tau=0} = X.$$

q.e.d.

With this in mind, we can prove the desired transversality result.

Lemma 7.4. *All moduli spaces $\overline{\mathcal{T}}(x_0, \vec{x})$ are manifolds of the expected dimension for a generic choice of universal perturbation data.*

Sketch of proof. If \vec{x} consists of a single element x_1 , this follows from the assumption that all functions $f_{i,j}$ are Morse-Smale. To prove the

general case, note that the choice of a perturbation datum associates a diffeomorphism of one of the manifolds $Q_{i,j}$ to each finite edge of a metric tree (T, g_T) ; the diffeomorphism is simply the flow of the associated perturbed gradient vector field. By picking finite segments in each incoming edge containing the support of the perturbation datum, we may also associate a diffeomorphism (the perturbed gradient flow on this segment) to a leaf of T . In particular, we define for each integer $0 \leq k \leq d$ a diffeomorphism ϕ_k obtained by composing the diffeomorphism attached to the k th leaf with the diffeomorphisms associated to the sequence of finite edges which form an arc from the k th leaf to the outgoing segment. The moduli space of gradient trees for a fixed metric tree is the inverse image of the diagonal in Q^{d+1} under a map

$$(7.12) \quad \text{ev}: W^s(x_0) \times W^u(x_1) \times \cdots \times W^u(x_d) \rightarrow Q^{d+1}$$

whose k th component ev_k is the composition of the inclusion of the given stable or unstable manifold in the appropriate factor with the diffeomorphism ϕ_k . A standard argument using Sard's theorem then reduces the regularity of the moduli space to the previous lemma.

q.e.d.

REMARK 7.5. Lemma 2.12 implies that we may discuss transversality of the evaluation map (7.12) to the diagonal even though the target space is not Hausdorff; the image of a point under consideration may only lie in the singular strata of Q if it is labeled by $Q_{1,1}$ or $Q_{2,2}$, in which case all the arguments take place in the appropriate manifold.

7.1.2. Transversality for shrubs. We briefly explain how to extend the techniques of the previous section to obtain moduli spaces of maps of shrubs which are smooth manifolds. The main difference are that the strata of $\overline{\mathcal{S}}_d$ are not manifolds with corners, so one may not extend an arbitrary smooth function from one skeleton to the next. The quickest way to resolve this problem is to consider the inclusion

$$(7.13) \quad \overline{\mathcal{S}}_d \rightarrow \overline{\mathcal{T}}_{2d} \equiv \overline{\mathcal{S}}_{2d}^{\text{ext}}$$

obtained by attaching two infinite edges to every incoming vertex of a shrub. We insist that this is a purely formal construction; the right-hand side should not be thought of as controlling gradient trees for Morse theory in this construction, and so we have given it a different name. It is clear that the natural stratification of $\overline{\mathcal{S}}_d$ agrees with the pull-back of the stratification of $\overline{\mathcal{T}}_{2d}$ under the inclusion.

Given a sequence \vec{I} , we define $2\vec{I}$ to be the sequence obtained by repeating every element but the last, allowing us to define the analogue of this inclusion in the presence of labels

$$(7.14) \quad \overline{\mathcal{S}}_{\vec{I}} \rightarrow \overline{\mathcal{T}}_{2\vec{I}} \equiv \overline{\mathcal{S}}_{2\vec{I}}^{\text{ext}}.$$

The reader may easily check the commutativity of the diagram,

$$(7.15) \quad \begin{array}{ccc} \overline{\mathcal{T}}_{\vec{R}} \times \overline{\mathcal{S}}_{\vec{I}^1} \times \dots \times \overline{\mathcal{S}}_{\vec{I}^r} & \longrightarrow & \overline{\mathcal{S}}_{\vec{I}} \\ \downarrow & & \downarrow \\ \overline{\mathcal{T}}_{\vec{R}} \times \overline{\mathcal{T}}_{2\vec{I}^1} \times \dots \times \overline{\mathcal{T}}_{2\vec{I}^r} & \longrightarrow & \overline{\mathcal{T}}_{2\vec{I}} \\ \downarrow \cong & & \downarrow \cong \\ \overline{\mathcal{T}}_{\vec{R}} \times \overline{\mathcal{S}}_{2\vec{I}^1}^{\text{ext}} \times \dots \times \overline{\mathcal{S}}_{2\vec{I}^r}^{\text{ext}} & \longrightarrow & \overline{\mathcal{S}}_{2\vec{I}}^{\text{ext}} \end{array}$$

with vertical arrows obtained by applying (7.14) and the horizontal ones by applying the map (3.5) in the first row, and r iterations of the map (2.26) in the second. Note that the map (2.26) can be interpreted in our new notation as the existence of a map

$$(7.16) \quad \overline{\mathcal{T}}_{\vec{I}_1} \times \overline{\mathcal{S}}_{\vec{I}_2}^{\text{ext}} \rightarrow \overline{\mathcal{S}}_{\vec{I}}^{\text{ext}}$$

where \vec{I} is obtained by removing one of the inputs of \vec{I}_1 and replacing it by the inputs of \vec{I}_2 .

Definition 7.6. A universal perturbation datum on $\overline{\mathcal{S}}^{\text{ext}}$ is a perturbation datum for every tree in $\overline{\mathcal{S}}_{\vec{I}}$ varying smoothly, and which is compatible with the perturbation datum $\mathbf{X}^{\mathcal{T}}$ used to define $\text{Morse}(Q_1, Q_2)$ and the maps (7.16).

A choice of perturbation data for each element of $\overline{\mathcal{S}}_{\vec{I}}$ is said to be **smooth** if it is the restriction of a universal perturbation datum on $\overline{\mathcal{S}}^{\text{ext}}$ under the map (7.14).

REMARK 7.7. Of course, a universal perturbation datum on $\overline{\mathcal{S}}^{\text{ext}}$ involves many choices which are discarded by restricting to $\overline{\mathcal{S}}_{\vec{I}}$. One might impose a universal perturbation datum only on neighborhoods of the image of $\overline{\mathcal{S}}_{\vec{I}}$ which are compatible with the operadic structure maps, but there seems to be no reason to desire this additional level of complexity.

The existence of a universal perturbation datum on $\overline{\mathcal{S}}^{\text{ext}}$ is proved by induction as we did in the construction of $\mathbf{X}^{\mathcal{T}}$. The induction step uses the fact that any smooth function on a subcomplex of the natural stratification of a manifold with corner may be extended to an additional cell.

The commutativity of the diagram (7.15) implies that the compatibility of a perturbation datum with respect to the structure maps of $\overline{\mathcal{S}}^{\text{ext}}$ implies compatibility of the restriction to $\overline{\mathcal{S}}$ with (3.5). Lemma 3.7 implies that the simplicial condition is realized in an open set of such perturbations. To conclude the desired transversality result, the key ingredient, as in the proof of Lemma 7.4 is an intersection-theoretic

interpretation of $\overline{\mathcal{S}}_{S, g_S}(x, \vec{\sigma})$ whenever $\vec{\sigma} = (\check{\sigma}_1, \dots, \check{\sigma}_d)$ is a composable sequence of cells with σ_k a cell of $\mathcal{Q}_{i_k, i_{k+1}}$, and x is a critical point of $f_{0,d}$. As we did for Stasheff trees in the proof of Lemma 7.4, this moduli space is the inverse image of the diagonal under a map

$$(7.17) \quad W^s(x) \times \check{\sigma}_1 \times \dots \times \check{\sigma}_d \rightarrow Q^{d+1}$$

whose k th factor is obtained by composing the inclusion of $\check{\sigma}_k$ with the composition of perturbed gradient flows associated to the edges of S lying on an arc from the k th incoming leaf to the outgoing edge. We conclude:

Lemma 7.8. *For generic universal perturbation data, all strata of $\overline{\mathcal{S}}(x, \vec{\sigma})$ are smooth manifolds of the expected dimension. q.e.d.*

7.1.3. Codimension 1 gluing for trees. We briefly describe how to use the approach of [12] in order to prove the desired gluing theorems for moduli spaces of gradient trees. The results of this section prove that the compactifications of moduli spaces of gradient trees are manifolds with boundary whenever they have virtual dimension 1. However, as we shall need a description of a neighborhood of the codimension 1 strata of moduli spaces of shrubs of arbitrary dimension, we shall discuss matters in that generality.

Consider a Morse-Smale function f on a smooth submanifold, and N_1 and N_2 smooth manifolds which are transverse to the gradient flow of f . The manifold of gradient flow lines of length R starting on N_1 and ending on N_2 will be denoted $\mathcal{M}^R(N_1, N_2)$. The union of these spaces admits a compactification $\overline{\mathcal{M}}(N_1, N_2)$ by allowing broken trajectories. We write

$$(7.18) \quad \mathcal{M}^\infty(N_1, x, N_2) \equiv W^s(x) \times W^u(x) \cap N_1 \times N_2 \subset Q \times Q$$

for broken gradient trajectories passing through a critical point x . Note that all strata of virtual codimension 1 in $\overline{\mathcal{M}}(N_1, N_2)$ are of this type. The following result is a minor generalization of Proposition 18.1.4 in [12]:

Lemma 7.9. *Each point $(\gamma_1, \gamma_2) \in \mathcal{M}^\infty(N_1, x, N_2)$ admits a neighborhood $U(N_1, x, N_2)$, and an embedding*

$$(7.19) \quad U(N_1, x, N_2) \times (R_0, +\infty] \rightarrow \overline{\mathcal{M}}(N_1, N_2)$$

for $R_0 \gg 0$, such that the image of $U(a_1, x, a_2) \times \{R\}$ consists of all gradient segments of length R which are sufficiently close to (γ_1, γ_2) . Moreover, the evaluation map

$$(7.20) \quad U(N_1, x, N_2) \times \{R\} \rightarrow Q \times Q$$

converges in the C^∞ topology as R converges to infinity. q.e.d.

Assuming that the perturbation data we have chosen are generic, this result implies that the moduli spaces $\overline{\mathcal{T}}(x_0, \vec{x})$ are manifolds with boundary, with boundary stratified by smooth manifolds of the expected dimension. For simplicity of notation (and as this is the only case that is needed for our arguments), we focus on a neighborhood of the strata of codimension 1. Such a stratum is a product

$$(7.21) \quad \mathcal{T}(x_0, \vec{x}^1) \times \mathcal{T}(y, \vec{x}^2)$$

with a fixed choice of an element $x_j^1 \in \vec{x}^1$ which equals y , and such that \vec{x} is the result of replacing x_j^1 by \vec{x}^2 in \vec{x}^1 . The outgoing and incoming edge with endpoint y are naturally isometric to $[0, +\infty)$ and $(-\infty, 0]$. We fix a constant K sufficiently large that the support of the perturbation data on these edges are contained, respectively, in $[0, K]$ and $[-K, 0]$, and write b_K for the image of $\pm K$ on this edge. We define evaluation maps

$$\begin{aligned} \mathcal{T}_{d_1} \times W^s(x_0) \times W^u(x_1^1) \times \cdots \times W^u(x_{j-1}^1) &\rightarrow Q^{d_1}, \\ &\times W^u(x_{j+1}^1) \times \cdots \times W^u(x_{d_1}^1) \\ \mathcal{T}_{d_2} \times W^u(x_2^1) \times \cdots \times W^u(x_2^{d_1}) &\rightarrow Q^{d_2} \end{aligned}$$

which are obtained in the second case by composing the inclusion of $W^u(x_2^j)$ into the j th factor with the flows associated with a descending arc from the j th leaf to the point b_k (the first case is a straightforward modification). We write $N^\infty(x_0, \vec{x}^1 - y)$ and $N^\infty(\vec{x}^2)$ for the inverse images of the diagonal under these maps. By definition, the stratum we are considering is given by

$$(7.22) \quad \mathcal{M}^\infty(N^\infty(x_0, \vec{x}^1 - y), y, N^\infty(\vec{x}^2)).$$

Our regularity assumption can be restated as the fact that this is a transverse fiber product.

Consider the ‘‘pregluing’’ map

$$(7.23) \quad (0, +\infty] \times \mathcal{T}_{d_1} \times \mathcal{T}_{d_2} \rightarrow \overline{\mathcal{T}}_d$$

obtained by replacing the bi-infinite edge in the source with a finite one whose length is given by the projection to the first factor in the source. Let us define

$$(7.24) \quad N^R(x_0, \vec{x}^1 - y) \text{ and } N^R(\vec{x}^2)$$

for $R \in [0, +\infty)$ in the same way as their analogue for $R = +\infty$, but using the perturbation data on $\overline{\mathcal{T}}_{\vec{r}}$ pulled back under the inclusion of $\{R\} \times \mathcal{T}_{\vec{r}_1} \times \mathcal{T}_{\vec{r}_2}$. A neighborhood of $\mathcal{T}(x_0, \vec{x}^1) \times \mathcal{T}(y, \vec{x}^2)$ in $\overline{\mathcal{T}}(y, \vec{x})$ consists of gradient trees whose source lies in the image of this pregluing map, and is therefore given by the union

$$(7.25) \quad \coprod_{0 \leq R \leq +\infty} \mathcal{M}^R(N^R(x_0, \vec{x}^1 - y), y, N^R(\vec{x}^2)).$$

The smoothness condition on our choice of perturbation data implies that the evaluation maps on $N^R(x_0, \vec{x}^1 - y)$ and $N^R(\vec{x}^2)$ converge on compact subsets, and in the C^∞ topology to the evaluation maps at $R = +\infty$. Together with Lemma 7.9, elementary differential topology implies:

Lemma 7.10. *Each point $(\psi_1, \psi_2) \in \mathcal{T}(x_0, \vec{x}^1) \times \mathcal{T}(y, \vec{x}^2)$ admits a neighborhood $U(\psi_1, \psi_2)$, and an embedding*

$$(7.26) \quad U(\psi_1, \psi_2) \times (R_0, +\infty] \rightarrow \overline{\mathcal{T}}(x, \vec{x})$$

for $R_0 \gg 0$, whose image is a neighborhood of (ψ_1, ψ_2) in the Gromov topology. *q.e.d.*

This result has a straightforward generalization to corners of higher codimension, but we shall not need that fact in this paper.

The same proof technique works for shrubs: Let ψ_0 be a rigid gradient tree in $\mathcal{T}(x_0, \vec{x})$, and $\vec{\psi} = (\psi_1, \dots, \psi_d)$ a sequence of shrub maps with outputs x_j . We write $(\vec{\sigma}^1, \dots, \vec{\sigma}^r)$ for the collections of inputs of all shrub maps ψ_j , and $\vec{\sigma}$ for the sequence obtained by concatenating the sequences $\vec{\sigma}^j$. By attaching ψ_j to the j th input of ψ_0 , we obtain a singular shrub

$$(7.27) \quad (\psi_0, \vec{\psi}) \in \overline{\mathcal{S}}(x_0, \vec{\sigma}).$$

Let $U(\psi_0, \vec{\psi})$ denote a sufficiently small neighborhood of this point in the boundary stratum of $\overline{\mathcal{S}}(x_0, \vec{\sigma})$ wherein it lies (note that this stratum has virtual codimension 1). Applying the discussion for trees to this situation, we conclude:

Lemma 7.11. *There exists an embedding*

$$(7.28) \quad U(\psi_0, \vec{\psi}) \times (R_0, +\infty] \rightarrow \overline{\mathcal{S}}(x_0, \vec{\sigma})$$

whose image is a neighborhood of $(\psi_0, \vec{\psi})$ in $\overline{\mathcal{S}}(x_0, \vec{\sigma})$. *q.e.d.*

7.2. Transversality and gluing for caps and mushrooms. The proofs of technical results from Section 5 are contained in this section. However, we omit the analytical parts of the proofs, so that we work with smooth functions throughout, without setting up the relevant Banach spaces.

7.2.1. Transversality for caps. In this section, we prove Lemma 5.11. One way to define the topology on $\mathcal{P}(\vec{p})$, is to introduce the space

$$C^\infty \mathcal{P}(\vec{p})$$

of all smooth maps $u: P \rightarrow M$ for some cap P whose boundary arcs are mapped to $(Q_{i_0}, \dots, Q_{i_d}, \mathcal{L}_q)$, where q is allowed to vary arbitrarily in Q , and whose incoming marked points are mapped to the Hamiltonian chords \vec{p} . Note that $\mathcal{P}(\vec{p})$ is simply the zero-level set of

$$(7.29) \quad (du - Y^{\vec{I}})^{0,1}$$

as a section of the vector bundle over $C^\infty\mathcal{P}(\vec{p})$ with fiber

$$C^\infty(P, \Omega^{0,1}(P) \otimes u^*TM)$$

at a map u with source P .

Proof of Lemma 5.11. The proof is a standard application of Sard-Smale to pseudo-holomorphic curve theory. We note that the assumption that all moduli spaces of holomorphic discs $\mathcal{R}(p_0, \vec{p})$ are regular implies that every stratum consisting of one cap and several discs will have the expected dimension as soon as the component corresponding to the cap is regular. We introduce the ingredients necessary to prove that this is the case for generic perturbation data:

Fix $q \in Q$, and let $u: P \rightarrow M$ be a smooth map. We define

$$T_u C^\infty\mathcal{P}(\vec{p})$$

to be the subspace of $C^\infty(P, u^*(TM))$ consisting of vector fields whose restriction to a boundary segment of P labeled by Q_i is valued in the tangent space Q_i , and whose restriction to the outgoing boundary segment projects to a constant vector in Q (note that the foliation \mathcal{L}_q defines a smooth map $M \rightarrow Q$). The projection map

$$(7.30) \quad T_u C^\infty\mathcal{P}(\vec{p}) \rightarrow T_q Q$$

is obviously surjective with kernel consisting of vector fields X which, on each arc point in the direction of the Lagrangian label. We write

$$(7.31) \quad C_{\text{res}}^\infty(P, u^*(TM))$$

for this space.

Let $\mathcal{T}^{\mathcal{P}}$ denote the space of infinitesimal deformations of the universal perturbation data $\mathbf{K}^{\mathcal{P}}$ and $\mathbf{J}^{\mathcal{P}}$, and let \mathcal{T}_P denote the analogous space for fixed P . By taking the direct sum of the linearized $\bar{\partial}$ operator and the projection above, we obtain a map

$$(7.32) \quad \mathcal{T}_P \oplus T_P\mathcal{P} \oplus T_u C^\infty\mathcal{P}(\vec{p}) \rightarrow C^\infty(P, \Omega^{0,1}(P) \otimes u^*TM) \oplus T_q Q.$$

Note that the surjectivity of this operator implies the desired result by the Sard-Smale theorem and the infinitesimal version of Lemma 5.9 which shows that $\mathcal{T}^{\mathcal{P}}$ surjects onto \mathcal{T}_P .

The surjectivity of this operator itself follows from the surjectivity of the linearized $\bar{\partial}$ operator restricted to the kernel of the projection to $T_P\mathcal{P} \oplus T_q Q$:

$$\mathcal{T}_P \oplus C_{\text{res}}^\infty(P, u^*(TM)) \rightarrow C^\infty(P, \Omega^{0,1}(P) \otimes u^*TM).$$

This fact, however, is a standard result in pseudo-holomorphic curve theory. For closed holomorphic curves, this appears in the textbook [13, Proposition 3.2.1] (the Lagrangian case can also be found, for example, in [18, Lemma 2.4]). q.e.d.

The next result is the key to our proof of the gluing theorem for mushrooms: Let u be an element of a moduli space of caps $\mathcal{P}(\vec{p}^1)$ of dimension k such that the evaluation map

$$(7.33) \quad \mathcal{P}(\vec{p}^1) \rightarrow Q$$

is an immersion near u ; let U denote a neighborhood in $\mathcal{P}(\vec{p}^1)$ where this property holds. Let v be a rigid element of a moduli space of discs $\mathcal{R}(p, \vec{p}^2)$ such that p agrees with the j th input of \vec{p}^1 . Denoting by \vec{p} the result of replacing this input in \vec{p}^1 with \vec{p}^2 , we write (u, v) for the singular cap in $\overline{\mathcal{P}}(\vec{p})$ obtained by attaching v to u . The next result is standard gluing theorem in pseudo-holomorphic curve theory:

Lemma 7.12. *There exists an embedding*

$$(7.34) \quad [0, 1) \times U \rightarrow \overline{\mathcal{P}}(\vec{p})$$

whose image is a neighborhood of (u, v) in $\overline{\mathcal{P}}(\vec{p})$ and whose restriction to $\{\epsilon\} \times U$ is smooth for any value of ϵ . Moreover, the composition of the evaluation map

$$(7.35) \quad \overline{\mathcal{P}}(\vec{p}) \rightarrow Q$$

with the inclusion of $\{\epsilon\} \times U$ converges in the C^1 topology to its value at $\epsilon = 0$.

Sketch of proof: The key ingredient in the construction of a gluing map is a choice of right inverse. Let q denote the image of u under the evaluation map, and fix a neighborhood of q with a diffeomorphism to a neighborhood of the origin in \mathbb{R}^n such that the image of U under the evaluation map is the intersection of this subset with \mathbb{R}^k .

We consider the codimension $n - k$ subspace

$$(7.36) \quad C_{\mathbb{R}^{n-k}}^\infty(P, u^*(TM)) \subset C_{\text{res}}^\infty(P, u^*(TM))$$

consisting of vector fields whose image under the projection map to TQ and the local diffeomorphism with \mathbb{R}^n lie in the copy of \mathbb{R}^{n-k} complementary to the image of the moduli space under evaluation. By construction, the restriction of the parametrized (perturbed) $\bar{\partial}$ operator to

$$(7.37) \quad T_P \mathcal{P} \oplus C_{\mathbb{R}^{n-k}}^\infty(P, u^*(TM))$$

is an isomorphism.

This choice of right inverse determines a gluing construction as follows: Since v is rigid, its $\bar{\partial}$ operator is an isomorphism. For every real number ϵ , one constructs a “preglued” map $u \#_\epsilon v$ whose source is a surface obtained by removing parts of the strip-like ends of the domains of u and v and attaching the complements. This preglued map comes equipped with a right inverse also obtained by gluing the right inverses on both sides. Applying the implicit function theorem to this right inverse, we conclude the existence of a unique vector field $\text{sol}_{u \#_\epsilon v}$ along

the pre-glued curve, in the image of the right inverse and of sufficiently small norm, whose image is an honest solution to the $\bar{\partial}$ problem being studied. The standard argument in holomorphic curve theory shows that the norm of $\text{sol}_{u\#\epsilon v}$ (in any reasonable space of functions) decays as ϵ goes to 0.

The same construction can be done for curves u' near u , and it is easy to see that the preglued curve varies smoothly with u and ϵ , from which C^0 convergence of the evaluation map follows. In order to prove C^1 convergence, we must therefore prove that the vector field $\text{sol}_{u\#\epsilon v}$ decays with ϵ in the C^1 topology as a function of u . The easiest way to do this is to use a right inverse at $u'\#\epsilon v$ obtained by parallel transport of the one constructed by pregluing at $u\#\epsilon v$. Decay in the C^1 topology follows from a standard version of the implicit function theorem applied to a chart for the relevant space of functions which is “centered at $u\#\epsilon v$ ” (see [6, Proposition 24] for the version of the implicit function theorem one should use, and [1, Proposition 5.4] for a detailed version of this argument in the setting of bubbling of holomorphic discs). q.e.d.

7.2.2. Transversality for mushrooms. In this section, we prove the existence of universal perturbation data for which the moduli spaces of mushrooms are smooth manifolds. We adopt the strategy used for shrubs and consider the disjoint union over sequences with $\sum_j d_j = d$:

$$(7.38) \quad \coprod \bar{\mathcal{S}}_{2r}^{\text{ext}} \times \bar{\mathcal{P}}_{d_1} \times \cdots \times \bar{\mathcal{P}}_{d_r}.$$

Recall that we constructed $\bar{\mathcal{C}}_d$ starting with a disjoint union of polyhedra each admitting a natural embedding:

$$(7.39) \quad \bar{\mathcal{S}}_{2r} \times \bar{\mathcal{P}}_{d_1} \times \cdots \times \bar{\mathcal{P}}_{d_r} \subset \bar{\mathcal{S}}_{2r}^{\text{ext}} \times \bar{\mathcal{P}}_{d_1} \times \cdots \times \bar{\mathcal{P}}_{d_r}.$$

The moduli space $\bar{\mathcal{C}}_d$ was defined as a quotient under an equivalence relation which naturally extends to the cells of the polyhedra appearing in the disjoint union (7.38). The cellular complex $\bar{\mathcal{C}}_d^{\text{ext}}$ obtained by taking the quotient under this relation is not a manifold (it doesn't even have pure dimension).

Let us consider the maps which make the moduli space of mushrooms into a bimodule over (two different models) of the Stasheff associahedron: A short verification shows that (4.4) is compatible with the inclusion (7.39), and a natural map

$$(7.40) \quad \bar{\mathcal{C}}_{d_1}^{\text{ext}} \times \bar{\mathcal{R}}_{d_2} \rightarrow \bar{\mathcal{C}}_{d_1+d_2-1}^{\text{ext}}$$

obtained by attaching the disc the one of the d_1 inputs of an element of $\bar{\mathcal{C}}_{d_1}^{\text{ext}}$. On the other hand, and as in Equation (7.16), the module map (4.3) can be recovered from the inclusion (7.39) and a repeated application of the map

$$(7.41) \quad \bar{\mathcal{T}}_{d_1} \times \bar{\mathcal{C}}_{d_2}^{\text{ext}} \rightarrow \bar{\mathcal{C}}_{d_1+d_2-1}^{\text{ext}}.$$

All these constructions have analogues for labels, so we shall consider a sequence \vec{I} of such. As all the components of (7.38) are smooth manifolds with corners, it makes sense to speak of smooth functions on $\overline{\mathcal{C}}_{\vec{I}}^{\text{ext}}$. We define a universal Q -parametrized perturbation datum in $\overline{\mathcal{C}}^{\text{ext}}$ to be a Q -parametrized perturbation datum on all elements of $\overline{\mathcal{C}}_{\vec{I}}^{\text{ext}}$ which is smooth on each component, invariant under the equivalence relation \sim , and compatible with the perturbation data $(\mathbf{X}^{\mathcal{T}}, \mathbf{K}^{\mathcal{R}}, \mathbf{J}^{\mathcal{R}})$ and the analogue of the maps (7.40) and (7.41) in the presence of labels.

Definition 7.13. *A family of perturbation data on $\overline{\mathcal{C}}_{\vec{I}}$ is **smooth** if it is the restriction of a universal perturbation datum on $\overline{\mathcal{C}}_{\vec{I}}^{\text{ext}}$.*

The analogue of Lemma 7.2 holds. As the proof is the same, we shall not repeat it. The only subtle point is ensuring that Condition (5.17) holds. However, the inductive construction starts by choosing perturbation data on the simplest space, which is the moduli of shrubs with one input. Definition 5.13 dictates how this initial choice should be made, and Condition (5.17) is compatible with the requirement that the perturbation data be chosen consistently.

Lemma 7.14. *The restriction map from the space of Q -parametrized perturbation data on $\overline{\mathcal{C}}_{\vec{I}}^{\text{ext}}$ to the space of perturbation data for a fixed mushroom is surjective. *q.e.d.**

The transversality result we desire now follows by applying the idea of Lemma 5.11. We start by studying interior strata of \mathcal{C}_d :

The strata of \mathcal{C}_d are labeled by the following data: a partition $d = d_1 + \dots + d_r$, an increasing sequence of integers $1 = r_0 < r_1 < \dots < r_e = r$, and a tree S with e inputs. Indeed, write \mathcal{D} for this data, and first consider the stratum of $\mathcal{S}_{\mathcal{D}} \subset \mathcal{S}_r$ consisting of trees homeomorphic to S for which two vertices v_i and v_j are collapsed to the same point whenever i and j lie between consecutive integers r_l and r_{l+1} , but are otherwise distinct. The stratum of \mathcal{C}_d consisting of mushrooms with r smooth caps with, respectively, (d_1, \dots, d_r) inputs and a stem in $\mathcal{S}_{\mathcal{D}}$ will be denoted by $\mathcal{C}_{\mathcal{D}}$; it is a product

$$(7.42) \quad \mathcal{S}_{\mathcal{D}} \times \mathcal{P}_{d_1} \times \dots \times \mathcal{P}_{d_r}.$$

Define

$$C^\infty \mathcal{P}_{(d_1, \dots, d_r)}(\vec{p}) \equiv C^\infty \mathcal{P}(\vec{p}_0^{d_1}) \times \dots \times C^\infty \mathcal{P}(\vec{p}_{d-d_r}^d),$$

to be the space of smooth r -tuples of maps from the disc to M satisfying the appropriate boundary conditions for the caps of a mushroom map in $\mathcal{C}(x_0, \vec{p})$ whose underlying mushroom has r smooth caps with, respectively, (d_1, \dots, d_r) inputs. Further, define

$$\mathcal{P}_{\mathcal{D}}(\vec{p}) \subset C^\infty \mathcal{P}_{(d_1, \dots, d_r)}(\vec{p}) \times \mathcal{C}_{\mathcal{D}}$$

to be the vanishing locus of the $\bar{\partial}$ equation (7.29) perturbed according to the chosen perturbation datum (which now depends on the modulus in $\mathcal{C}_{\mathcal{D}}$ rather than depending only on the corresponding moduli space of caps as explained in Remark 5.20).

We have a map

$$(7.43) \quad \mathcal{P}_{\mathcal{D}}(\vec{p}) \rightarrow \mathcal{C}_{\mathcal{D}} \times Q^r$$

which, on the second factor, agrees with the projection of the “outgoing arc” of each cap.

Let us now consider the stem. Generalizing Equation (7.17), and writing $\vec{Q} = (Q_{i_0}, \dots, Q_{i_r})$ we define

$$(7.44) \quad \mathcal{S}_{\mathcal{D}}(x_0, \vec{Q})$$

to be the inverse image of the diagonal under a map

$$(7.45) \quad \mathcal{C}_{\mathcal{D}} \times W^s(x_0) \times Q_{i_0, i_0^2} \times \cdots \times Q_{i_r, i_d} \rightarrow Q^{r+1},$$

whose factors are a composition of the natural inclusion with a diffeomorphism of the target obtained by integrating the flow of a descending arc of S as in Equation (7.17) (the flow depends on the perturbation data, which are parametrized by the first factor). Here, as in Equation (5.15), the sequence $\{i_j^k\}_{j=0}^{d_k}$ records the boundary conditions on $\mathcal{P}(\vec{p}^k)$. Observe that the moduli space of shrubs with a fixed source, output x_0 , and arbitrary inputs is naturally an open submanifold of the ascending manifold of x_0 . The reader should also keep in mind that while there are r evaluation maps, we are potentially considering a tree in which an arbitrary number of these maps necessarily agree because the corresponding incoming vertices lie on a subset of the shrub which is collapsed to a point.

The next result is simply a reinterpretation of the definition of a mushroom map (compare also Equation (5.15)):

Lemma 7.15. *The subset of $\mathcal{C}(x_0, \vec{p})$ consisting of maps whose source is a mushroom in $\mathcal{C}_{\mathcal{D}}$ is naturally diffeomorphic to the fiber product*

$$(7.46) \quad \mathcal{C}_{\mathcal{D}}(x_0, \vec{p}) = \mathcal{S}_{\mathcal{D}}(x_0; \vec{Q}) \times_{\mathcal{C}_{\mathcal{D}} \times Q^r} \mathcal{P}_{\mathcal{D}}(\vec{p}).$$

q.e.d.

We complete this section by proving a Lemma from Section 5.3.

Proof of Lemma 5.22. Since every component of a singular mushroom map consists of mushroom maps, together with discs perturbed with respect to the datum $(\mathbf{K}^{\mathcal{R}}, \mathbf{J}^{\mathcal{R}})$, as well as gradient Stasheff trees perturbed with respect to $\mathbf{X}^{\mathcal{T}}$, the fact that we have chosen regular perturbation data to define $\text{Morse}(Q_1, Q_2)$ and $\text{Fuk}(Q_1, Q_2)$ implies that Lemma 5.22 would follow once we show that all strata of $\mathcal{C}(x_0, \vec{p})$ have the expected dimension. The previous lemma reduces this to showing

that for a generic choice of perturbation data, the fiber product (7.46) is transverse. Of course, this will follow from an application of Sard-Smale. To simplify matters, we observe that the forgetful map

$$\mathcal{S}_{\mathcal{D}}(x_0; \vec{Q}) \rightarrow \mathcal{S}_{\mathcal{D}}$$

is a submersion. In particular, our result would follow once we show that the evaluation map

$$(7.47) \quad \mathcal{P}_{\mathcal{D}}(\vec{p}) \rightarrow \mathcal{P}_{d_1} \times \cdots \times \mathcal{P}_{d_r} \times Q^r$$

can be made transverse to an arbitrary submanifold in the target by changing the perturbation datum.

Fix a mushroom map Ψ with source C , and let $\mathcal{T}_{\vec{P}(C)}$ denote the space of infinitesimal deformations of perturbation data for the caps of C . We introduce the notation

$$(7.48) \quad C_{\text{res}}^{\infty}(\vec{P}, \Psi^*(TM)) = \bigoplus_{1 \leq j \leq r} C_{\text{res}}^{\infty}(P_j(C), \Psi_j^*(TM)).$$

Using the surjectivity of the restriction map from the space of universal perturbation data to the space of perturbation data for the caps of C , we have reduced the problem to showing that the evaluation map at the level of tangent spaces

$$(7.49) \quad \mathcal{T}_{\vec{P}(C)} \times C_{\text{res}}^{\infty}(\vec{P}, \Psi^*(TM)) \rightarrow TQ^r$$

is surjective. Splitting the source and the target as a direct sum of r factors, we find that this follows immediately from the proof of Lemma 5.11 given in the previous section.

To prove transversality even in the base case of mushrooms with one input, it is important to note that surjectivity holds even if we are allowing only perturbations of the almost complex structure (but which are of course allowed to vary along the caps, this is exactly the setting in which transversality is proved in [13]). q.e.d.

7.2.3. Mushroom gluing: the rigid case. Let us now consider the situation considered in Proposition 5.23: a critical point x_0 and a sequence of time-1 Hamiltonian chords \vec{p} such that the moduli space

$$(7.50) \quad \mathcal{C}(x_0, \vec{p})$$

has virtual dimension 1. By Lemma 5.22, which is proved in the previous section, all strata of $\overline{\mathcal{C}}(x_0, \vec{p})$ are empty except those of virtual dimension 0 and 1. In order to prove that $\overline{\mathcal{C}}(x_0, \vec{p})$ is a 1-dimensional manifold with boundary, we must prove that points in the 0-dimensional strata have neighborhoods, in the Gromov topology, which are homeomorphic to an interval.

There are four distinct geometric situations to consider.

- 1) ψ is a rigid gradient tree in $\mathcal{T}(x_0, \vec{x})$, and $\vec{\Psi} = (\Psi_1, \dots, \Psi_k)$ is a sequence of rigid mushrooms with outputs x_j . We write $(\vec{p}^1, \dots, \vec{p}^r)$ for the collections of inputs of all mushrooms Ψ_j , and \vec{p} for the sequence obtained by concatenating the sequences \vec{p}^j . The pair $(\psi, \vec{\Psi})$ defines an element of $\overline{\mathcal{C}}(x_0, \vec{p})$ obtained geometrically by attaching Ψ_j to the j th input of ψ . Consider the stratum $\overline{\mathcal{C}}_{\mathcal{D}}$ determined by the datum of $d_j = |\vec{p}^j|$ and a tree S with exactly r inputs. If the perturbation data on the caps were independent of the stem, the result would follow from Lemma 7.11, and the observation that the moduli space we are studying has an alternative description as

$$(7.51) \quad \overline{\mathcal{S}}(x_0, \mathcal{P}(\vec{p}^1, \dots, \vec{p}^r))$$

where $\mathcal{P}(\vec{p}^1, \dots, \vec{p}^r)$ is the product of the moduli spaces $\mathcal{P}(\vec{p}^j)$, each of which defines a chain on Q by the evaluation map at the outgoing segment. Using only Lemma 7.11 as analytic input one can prove the general case by imitating the strategy for passing from Equation (5.15) to Lemma 7.15 (i.e., work fiberwise over $\overline{\mathcal{C}}_{\mathcal{D}}$).

- 2) v is a rigid holomorphic curve in $\mathcal{R}(p_0, \vec{p}^2)$, and Ψ a rigid mushroom in $\overline{\mathcal{C}}(x_0, \vec{p}^1)$, such that p_0 agrees with the k th input in \vec{p}^1 . In this case, we focus attention on the cap with inputs $\vec{p}^{1,j}$ wherein this k th input lies. Writing Ψ_j for the map of caps along which we are attaching v , we assume again for simplicity that the perturbation datum on that cap is independent of the remaining caps and the stem, and use Lemma 7.12 to find a neighborhood of (Ψ_j, v) in $\overline{\mathcal{P}}(\vec{p}^{1,j} \cup_{p_0} \vec{p}^2)$ foliated by smooth codimension 1 submanifolds which are the images of the boundary for gluing parameter ϵ . Note that the immersive assumption in Lemma 7.12 is implied by the rigidity and regularity requirements on Ψ . In addition, Lemma 7.15 gives a description of $\overline{\mathcal{C}}(x_0, \vec{p} \#_{p_0} \vec{p}^2)$ as a fiber product one of whose factors is $\overline{\mathcal{P}}(\vec{p}^1 \#_{p_0} \vec{p}^2)$. The assumption that $\overline{\mathcal{C}}(x_0, \vec{p})$ is rigid and the C^1 convergence of the evaluation maps to Q on the image of the gluing map implies that there is a unique solution for each such sufficiently small parameter. The general case can be proved by working fiberwise over $\mathcal{C}_{\mathcal{D}}$, or, alternatively, by noticing that the smoothness in the choice of perturbation data implies that there is a neighborhood of the singular mushroom we are considering where all moduli spaces together with their evaluation maps are perturbed by C^∞ small amounts from moduli spaces with perturbation data independent of the stem. After such a perturbation, a unique ray converging to the boundary persists.
- 3) Ψ is a mushroom with r caps, such that two successive incoming vertices (say, $k-1$ and k) of the stem are collapsed to one point, which is moreover rigid among such mushrooms. In particular, the

boundary conditions on the outgoing arcs of the two caps $P_{k-1}(C)$ and $P_k(C)$ agree. In this case, there is a unique 1-parameter family of mushrooms Ψ^t , whose domains have stems for which the distance between the incoming vertex labeled k and $k+1$ does not vanish, and which converge to Ψ . The argument given in Lemma 3.10 proves this.

- 4) In the same situation as before, we glue $\Psi_{P_{k-1}}$ and Ψ_{P_k} to obtain a cap whose inputs are given by concatenating those for $\Psi_{P_{k-1}}$ and Ψ_{P_k} . Lemma 7.12 again applies here to show that this moduli space of caps is foliated by the images of the gluing map for sufficiently small gluing parameter ϵ , and that the evaluation map converges in the C^1 topology as ϵ goes to 0. The argument of Case (2) yields the desired conclusion.

8. Orientations and signs

8.1. The differential in Morse and Floer theory. We begin by describing the differential in the Morse complex $CM^*(f_{i,j})$. We define $\mathcal{T}(x_0, x_1)$ to be the fiber product of $W^s(x_0)$ and $W^u(x_1)$ over their common inclusion in $Q_{i,j}$, yielding a natural short exact sequence of vector spaces

$$(8.1) \quad 0 \rightarrow T\mathcal{T}(x_0, x_1) \rightarrow T(W^s(x_0) \times W^u(x_1)) \rightarrow TQ_{i,j} \rightarrow 0,$$

and hence an isomorphism of top exterior powers

$$(8.2) \quad \lambda(\mathcal{T}(x_0, x_1)) \otimes \lambda(Q_{i,j}) \cong \lambda(W^s(x_0)) \otimes \lambda(W^u(x_1)).$$

Whenever $\deg(x_0) = \deg(x_1) + 1$, the intersection between $W^s(x_0)$ and $W^u(x_1)$ is a segment, mapping to \mathbb{R} immersively under $f_{i,j}$, so the tangent space of

$$(8.3) \quad \psi \in \mathcal{T}(x_0, x_1)$$

is canonically oriented, with our conventions dictating that a rigid gradient flow line is given the opposite orientation of the one induced by projection to \mathbb{R} (i.e. the orientation we pick identifies x_1 with $-\infty$). Using the decomposition of the tangent space of $Q_{i,j}$ at x_1 , we have an isomorphism

$$(8.4) \quad \lambda(Q_{i,j}) \cong \lambda(W^s(x_1)) \otimes \lambda(W^u(x_1)),$$

which, together with the map in (8.2) determines an isomorphism

$$(8.5) \quad \lambda(W^s(x_1)) \cong \lambda(W^s(x_0)),$$

and hence an induced map on orientation lines

$$(8.6) \quad \mu_\psi^M: |o_{x_1}| \rightarrow |o_{x_0}|.$$




Figure 10

The differential is defined to be the sum of all such terms multiplied by a global sign:

$$(8.7) \quad \mu_1^M([x_1]) = \sum_{\substack{\deg(x_0)=\deg(x_1)+1 \\ \psi \in \mathcal{T}(x_0, x_1)}} (-1)^n \mu_\psi^M([x_0]).$$

8.2. Orienting abstract moduli spaces. As noted in our discussion of the Fukaya and Morse categories, the construction of A_∞ categories from counts of rigid discs or trees requires first a choice of orientations on the abstract moduli spaces controlling these categories.

8.2.1. Stasheff trees. Let T be a trivalent Stasheff tree with d inputs. Given an integer k , let A_k denote the descending arc starting at the k th incoming leaf, and let b_k denote the vertex of T where A_k and A_{k-1} meet. We shall label the edges of this trivalent vertex e_k^ℓ , e_k^r and e_k^d as in Figure 10.

We order the (finite) edges of T using the following inductive procedure:

- 1) Set T_2 to be the union of all the external edges of T .
- 2) If e_k^r lies in T_k , then define $e_k = e_k^d$. Otherwise, set $e_k = e_k^r$.
- 3) Set T_k to be the union of T_{k-1} with e_k .

Lemma 8.1. *If $j \neq k$, $e_j \neq e_k$.*

Proof. We must prove that it is impossible for both e_k^d and e_k^r to lie in T_{k-1} . The edge e_k^d lies in T_{k-1} only if it is the outgoing external edge. In that case, a simple check shows that e_k^r cannot lie in T_{k-1} . q.e.d.

We define the orientation on \mathcal{T}_d to be given by

$$(8.8) \quad (-1)^{r(T)} dt_{e_3} \wedge \cdots \wedge dt_{e_d}$$

where t_e is the length of e , and $r(T)$ is the number of integers between 3 and d for which $e_k = e_k^r$.

Lemma 8.2. *The orientation given by formula (8.8) is independent of the topological type of T .*

Proof. It suffices to check that this is the case in a neighborhood of the codimension 1 walls. As these consist of trees with a unique

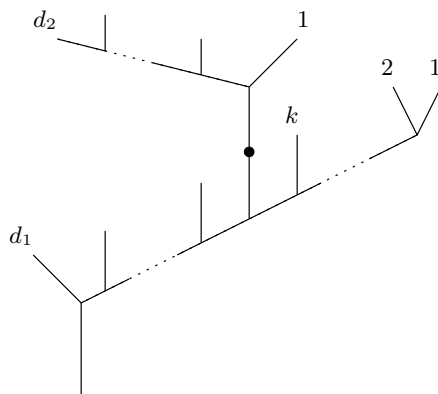


Figure 11

4-valent vertex, we start with a trivalent tree, and pass through the wall by collapsing an edge e . There is a natural way to identify the edges on trivalent trees on either side of the wall, and the reader may easily check that, under this identification, the edges e_j are unchanged whenever b_k is not one of the endpoints of e . There are two cases to consider, depending on whether e_k agrees with e_k^r , whose verification is left to the reader. q.e.d.

We now consider the boundary stratum of $\overline{\mathcal{T}}_d$ which is the image of the map

$$(8.9) \quad \overline{\mathcal{T}}_{d_1} \times \overline{\mathcal{T}}_{d_2} \rightarrow \overline{\mathcal{T}}_d$$

given by attaching a tree with d_2 inputs to the $(k+1)$ st incoming leaf of a tree with d_1 inputs (here, $d_1 + d_2 = d + 1$). Since we know that our choices of orientations are consistent throughout the moduli space, it suffices to compare these orientations in a special case. The author found the case shown in Figure 11 to be particularly convenient. The resulting computation implies:

Lemma 8.3. *The product orientation on*

$$(8.10) \quad (-1, 0] \times \overline{\mathcal{T}}_{d_1} \times \overline{\mathcal{T}}_{d_2} \hookrightarrow \overline{\mathcal{T}}_d$$

differs from the boundary orientation by a sign of

$$(8.11) \quad (d_1 - k)d_2 + d_2 + k.$$

q.e.d.

8.2.2. Shrubs. The moduli space of shrubs can be most easily oriented using the map

$$(8.12) \quad \mathcal{S}_d \rightarrow \mathcal{T}_d$$

obtained by replacing every incoming leaf by one which has infinite length. The fiber is isomorphic to a half-infinite line, which we trivialize




Figure 12

using the distance from the outgoing edge to an incoming leaf. We write ℓ for this function, and use the orientation

$$(8.13) \quad (-1)^{r(T)} d\ell \wedge dt_{e_3} \wedge \cdots \wedge dt_{e_d}$$

for shrubs of homeomorphism type T (note that the homeomorphism type is preserved by projection to \mathcal{T}_d). Lemma 8.2 implies immediately that this formula gives orientations on the top strata of \mathcal{S}_d compatible along their common boundary facets. It remains to understand the inductive behavior at the boundary of $\overline{\mathcal{S}}_d$:

Lemma 8.4. *The product orientation on*

$$(8.14) \quad (-1, 0] \times \overline{\mathcal{T}}_r \times \overline{\mathcal{S}}_{d_1} \times \cdots \times \overline{\mathcal{S}}_{d_r} \rightarrow \overline{\mathcal{S}}_d$$

differs from the boundary orientation by a sign of

$$(8.15) \quad 1 + \sum_{k=1}^r (r - k)(d_k + 1)$$

while the orientation on the stratum

$$(8.16) \quad (-1, 0] \times \overline{\mathcal{S}}_{d-1} \rightarrow \overline{\mathcal{S}}_d$$

whose image consists of shrubs whose k th and $(k - 1)$ st inputs are collapsed differs by a sign of

$$(8.17) \quad k + 1$$

from its orientation as a subset of the boundary.

Proof. Given the fact that the orientations are consistent throughout the moduli space, it suffices to compute the difference in sign orientation for one topological type of trees. The case (8.16) is essentially trivial, while that of (8.14) can be conveniently checked using the singular shrub appearing in Figure 12. q.e.d.

8.3. Orienting moduli spaces of shrubs. We shall fix the following conventions: whenever a space X is given as a transverse fiber product of X_1 and X_2 over X_0 , the short exact sequence

$$(8.18) \quad 0 \rightarrow TX \rightarrow TX_1 \oplus TX_2 \rightarrow TX_0 \rightarrow 0$$

determines a fixed isomorphism

$$(8.19) \quad \lambda(X) \otimes \lambda(X_0) \cong \lambda(X_1) \otimes \lambda(X_2).$$

For example, the isomorphism of Equation (2.34) comes from the choices

$$(8.20) \quad \begin{aligned} X &= \mathcal{T}(x_0, \vec{x}), & X_0 &= Q^{d+1}, \\ X_1 &= Q, & X_2 &= \mathcal{T}_d \times W^s(x_0) \times W^u(x_1) \times \cdots \times W^u(x_d), \end{aligned}$$

while Equation (3.13) comes from the choices

$$(8.21) \quad \begin{aligned} X &= \mathcal{S}(x_0, \vec{\sigma}), & X_0 &= Q^{d+1}, \\ X_1 &= Q, & X_2 &= \mathcal{T}_d \times W^s(x_0) \times \check{\sigma}_1 \times \cdots \times \check{\sigma}_d. \end{aligned}$$

We first use this fixed choice in order to prove the validity of the signs in Proposition 3.10. More precisely, we shall prove the validity of the signs in the right-hand side, leaving the left hand side to the reader. For clarity of the argument, we introduce a change of sign of

$$(8.22) \quad (n+1)(r + \dagger(\vec{\sigma}))$$

which we use to change the orientation of all moduli spaces of shrubs with inputs $\vec{\sigma}$ and dimension 1.

The starting point is an element of $\overline{\mathcal{S}}(x_0, \vec{\sigma})$ which consists of a rigid gradient tree in $\mathcal{T}(x_0, \vec{x})$ with $\vec{x} = (x_1, \dots, x_r)$, and a collection of r rigid shrubs in $\mathcal{S}(x_r, \vec{\sigma}^r)$, where $(\vec{\sigma}^1, \dots, \vec{\sigma}^r)$ is a partition of $\vec{\sigma}$. The orientation of the rigid tree is given by Equation (2.34) which can be simplified to

$$\lambda(Q^r) \cong \lambda(\mathcal{T}_r) \otimes \lambda(W^s(x_0)) \otimes \lambda(W^u(\vec{x}))$$

by using the splitting of the inclusion of the tangent space of the $(r+1)$ -fold diagonal whose kernel are the last r factors. Next, we decompose the k th copy of $\lambda(Q^r)$ on the left as a tensor product

$$\lambda(W^s(x_k)) \otimes \lambda(W^u(x_k)),$$

and rearrange the terms in the left hand side as

$$\lambda(W^s(\vec{x})) \otimes \lambda(W^u(\vec{x})),$$

where the first factor is the tensor product of $\lambda(W^s(x_k))$ in their usual order, introducing a Koszul sign of

$$(8.23) \quad \sum_{k=1}^r \deg(x_k) \left(n(k-1) + \sum_{j=1}^{k-1} \deg(x_j) \right).$$

Cancelling $\lambda(W^u(\vec{x}))$ from both sides of the equation, we arrive at the isomorphism

$$(8.24) \quad \lambda(W^s(\vec{x})) \cong \lambda(\mathcal{T}_r) \otimes \lambda(W^s(x_0)).$$

Rigid shrubs are oriented via Equation (2.34), which can be simplified (using the same splitting of the inclusion of the tangent space of the diagonal) as

$$\lambda(Q^{d_k}) \cong \lambda(\mathcal{S}_{d_k}) \otimes \lambda(W^s(x_k)) \otimes \lambda(\vec{\sigma}^k).$$

Introducing a sign of

$$(8.25) \quad \sum_{k=1}^r \deg(x_k)(d_k + 1),$$

we obtain an isomorphism

$$(8.26) \quad \lambda^{-1}(W^s(x_k)) \otimes \lambda(Q^{d_k}) \cong \lambda(\mathcal{S}_{d_k}) \otimes \lambda(\vec{\sigma}^k).$$

The product orientation of the singular shrub we are considering is given by taking the tensor product of Equation (8.24) with a copy of Equation (8.26) for each shrub:

$$\begin{aligned} \lambda(W^s(\vec{x})) \otimes \lambda^{-1}(W^s(x_1)) \otimes \lambda(Q^{d_1}) \otimes \cdots \otimes \lambda^{-1}(W^s(x_r)) \otimes \lambda(Q^{d_r}) &\cong \\ \lambda(\mathcal{T}_r) \otimes \lambda(W^s(x_0)) \otimes \lambda(\mathcal{S}_{d_1}) \otimes \lambda(\vec{\sigma}^1) \otimes \cdots \otimes \lambda(\mathcal{S}_{d_r}) \otimes \lambda(\vec{\sigma}^r). \end{aligned}$$

We shall compare this with the induced orientation as a boundary. On the left-hand side, we permute factors so that we arrive at

$$\left(\bigotimes_{k=1}^r \lambda^{-1}(W^s(x_k)) \otimes \lambda(W^s(x_k)) \right) \otimes \lambda(Q^d),$$

which introduces a sign of

$$(8.27) \quad \sum_{k=1}^r \left(\deg(x_k) \left(n(d_1 + \cdots + d_{k-1}) + \sum_{j=k}^r \deg(x_j) \right) \right).$$

On the right-hand side, we rearrange the terms so that they appear in the order

$$\lambda(\mathcal{T}_r) \otimes \lambda(\vec{\mathcal{S}}_{d_1}) \otimes \cdots \otimes \lambda(\vec{\mathcal{S}}_{d_r}) \otimes \lambda(W^s(x_0)) \otimes \lambda(\vec{\sigma}).$$

The Koszul signs amount to

$$\sum_{k=1}^r (d_k + 1) \left(\deg(x_0) + \sum_{j=1}^{k-1} (\deg(\vec{\sigma}^j) + nd_j) \right),$$

which we prefer to rewrite as

$$(8.28) \quad (r+d) \deg(x_0) + \sum_{j=1}^{k-1} n d_j (d_k + 1) + \sum_{k=1}^r \deg(\vec{\sigma}^k) \left(r - k - 1 + \sum_{j=k+1}^r d_j \right).$$

The difference between the boundary and product orientations is given by the sum of Equations (8.23), (8.25), (8.27), and (8.28). The first three give

$$\sum_{k=1}^r \deg(x_k) \left((k+1)(n+1) + k + d_k + n(d_1 + \cdots + d_{k-1}) + r + \deg(x_0) \right).$$

The term $\deg(x_k)(k+1)(n+1)$ cancels with a term appearing in the definition of the Morse category, which introduces an additional sign of $(n+1)r$. This can be seen by writing

$$\deg(x_0) + \sum_{k=1}^r k \deg(x_k) = r + \sum_{k=1}^r (k-1) \deg(x_k)$$

using the fact that the gradient tree with inputs x_k and output x_0 is rigid. Incorporating these signs coming from the Morse category yields

$$(n+1)r + \sum_{k=1}^r (\deg(\vec{\sigma}^k) + d_k + 1) (1 + k + d_k + n(d_1 + \cdots + d_{k-1})),$$

and the sum with (8.28) equals

$$(r+d) \deg(x_0) + (n+1)r + \sum_{k=1}^r (d_k + 1)(d_k + 1 + k) + \sum_{k=1}^r \deg(\vec{\sigma}^k) \left(r + d + (n+1) \sum_{j=1}^{k-1} d_j \right).$$

There are additional signs coming from the definition of the functor \mathcal{F} (see Equation (3.17)), and the change in orientations of 1-dimensional moduli spaces given by Equation (8.22). Together, these terms contribute

$$(8.29) \quad (n+1) \left(r + \sum_{k=1}^r \left(\deg(\sigma_k) \sum_{j=1}^{k-1} d_k \right) \right),$$

after which we use the fact that $\sum \deg(\vec{\sigma}) = \deg(x_0) + d$ to simplify the total remaining sign to

$$\begin{aligned} (r+d)d + \sum_{k=1}^r (d_k+1)(1+d_k+k) &= (r+1)(d+r) + \sum_{k=1}^r (d_k+1)(1+k) \\ &= \sum_{k=1}^r (d_k+1)(r+k). \end{aligned}$$

Adding the sign coming from Equation (8.15), we conclude that the boundary and product orientations are exactly opposite with the convention we have chosen.

8.4. Orientations for mushroom maps. Throughout this section, we shall work with the simplified setup in which moduli spaces of mushroom maps are given by the fiber product in (5.15). More precisely, the moduli space of mushrooms $\mathcal{C}(x_0, \vec{p})$ is stratified with top strata labeled by a choice of a partition of \vec{p} into sequences $(\vec{p}^1, \dots, \vec{p}^r)$. The stratum determined by this data is the moduli space of shrubs with r inputs lying on the images of the projection maps from the moduli spaces of caps $\{\mathcal{P}(\vec{p}^k)\}_{k=1}^r$, and is therefore given by a fiber product as in the previous section with

$$(8.30) \quad \begin{aligned} X_0 &= Q^{r+1}, & X &= \mathcal{S}(x_0, \mathcal{P}(\vec{p}^1), \dots, \mathcal{P}(\vec{p}^r)), \\ X_1 &= Q, & X_2 &= \mathcal{S}_r \times W^s(x_0) \times \mathcal{P}(\vec{p}^1) \times \dots \times \mathcal{P}(\vec{p}^r). \end{aligned}$$

Every such mushroom Ψ therefore determines an isomorphism

$$(8.31) \quad \begin{aligned} \lambda(\mathcal{C}(x_0, \vec{p})) \otimes \lambda(Q^{r+1}) &\cong \\ &\lambda(Q) \otimes \lambda(\mathcal{S}_r) \otimes \lambda(W^s(x_0)) \times \lambda(\mathcal{P}(\vec{p}^1)) \times \dots \times \lambda(\mathcal{P}(\vec{p}^r)). \end{aligned}$$

The description of $\lambda(\mathcal{P}(p_1))$ given in Equation (5.20) extends to an isomorphism

$$(8.32) \quad \lambda(\mathcal{P}(\vec{p}^k)) \cong \lambda(\mathcal{P}_{d_k}) \otimes \mathfrak{o}_{p_1^k}^\vee \otimes \dots \otimes \mathfrak{o}_{p_{d_k}^k}^\vee \otimes \lambda(Q).$$

When Ψ is rigid, our choice of orientations on \mathcal{S}_r , on \mathcal{R}_{d_k+1} (using the conventions from [17]) and the identification $\mathcal{P}_{d_k} \cong \mathcal{R}_{d_k+1}$ fixed in Remark 4.2, therefore determine a unique homomorphism

$$(8.33) \quad \mathfrak{o}_{p_1} \otimes \dots \otimes \mathfrak{o}_{p_d} \rightarrow \mathfrak{o}_{x_0}.$$

We write \mathcal{F}^Ψ for the induced map on orientation bundles. In order to get the correct sign, we first introduce the notation

$$(8.34) \quad \dagger(\vec{P}) \equiv \sum_{j=1}^r k \deg(\mathcal{P}(\vec{p}^k)) = \sum_{j=1}^k \left(1 - d_k + \sum_{\ell=1}^{d_k} \deg(\vec{p}^{\ell k}) \right),$$

where $\deg(\vec{p}^k)$ is the sum of the degrees of the elements of \vec{p}^k . The sign used to define \mathcal{G} is given by

$$(8.35) \quad \ddagger(\Psi) = (n+1) \ddagger(\vec{\mathcal{P}}) + \sum_{k=1}^r \ddagger(\vec{p}^k).$$

Having fixed the signs in the functor \mathcal{F} , we have the necessary ingredients to prove Proposition 4.2. The proof considers the four types of boundary strata discussed in Section 7.2.3 and shows that each such stratum contributes with the correct sign to Equation (5.32), after taking into account the sign (8.35), as well as the signs that go into the construction of $\text{Fuk}(Q_1, Q_2)$ and $\text{Morse}(Q_1, Q_2)$. More precisely, it is enough to show that, after taking these artificial signs into account, the difference between the boundary and product orientation on each stratum of $\vec{\mathcal{C}}(x_0, \vec{p})$ is given by the sign claimed in the A_∞ equation (3.30), up to a universal constant, independent of the stratum (but that is allowed to depend on \vec{p}).

The proof is a tedious computation, with the cases (1) and (2) from Section 7.2.3 following, respectively, from the computations performed in [17] and the previous section. Indeed, whenever a rigid disc bubbles off the k th cap, the only affected term in Equation (8.35) is $\ddagger(\vec{p}^k)$. This is exactly the sign convention for holomorphic discs we have used in constructing the Fukaya category (see [17, Equation (12.24)]). Similarly, when the stem breaks into a gradient tree together with several shrubs, only the term $(n+1) \ddagger(\vec{\mathcal{P}})$ changes. The integer $\deg(\mathcal{P}(\vec{p}^k))$ agrees with the codimension of $\mathcal{P}(\vec{p}^k)$ as a chain on Q , so these are exactly the conventions used in the previous section while orienting moduli space of shrubs.

Therefore, we shall only discuss cases (3) and (4), which correspond to terms which are supposed to cancel pairwise. In terms of orientations, we must therefore show that, whenever given a singular mushroom map lying in such an “internal” boundary, there is no difference between the orientations obtained by considering it as a boundary stratum of the two components described in cases (3) and (4). As in the previous section, it is convenient to change the orientation of the 1-dimensional moduli spaces as well. The sign we use is given by

$$(8.36) \quad (n+1) \left(r + \ddagger(\vec{\mathcal{P}}) \right) + \sum_{k=1}^r \left(1 + \ddagger(\vec{p}^k) + d_k \deg(\vec{p}^k) \right).$$

We will start with case (3), for which the orientation is induced by the isomorphism of Equation (8.31). For simplicity, we begin by choosing a splitting of the inclusion of the tangent space $(r+1)$ -diagonal of Q , corresponding to the inclusion of all but the first factor. Introducing a

sign of n , we can simplify to an isomorphism

$$(8.37) \quad \lambda(\mathcal{C}(x_0, \vec{p})) \otimes \lambda(Q^r) \cong \lambda(\mathcal{S}_r) \otimes \lambda(W^s(x_0)) \otimes \lambda(\mathcal{P}(\vec{p}^1)) \otimes \cdots \otimes \lambda(\mathcal{P}(\vec{p}^r)).$$

The next step is to use the isomorphism (8.32) to decompose $\lambda(\mathcal{P}(\vec{p}^{k-1}))$ into its constituent factors, then cancel the resulting copy of $\lambda(Q)$ with its matching factor on the right-hand side (this is the $(k-1)$ st factor starting at the left). The resulting Koszul sign is

$$(8.38) \quad n(k+1) + n \left(r + 1 + \deg(x_0) + n(k-2) + \sum_{j=1}^{k-1} \deg(\mathcal{P}(\vec{p}^j)) \right) \\ = n \left(r + \deg(x_0) + \sum_{j=1}^{k-1} (1 + d_j + \deg(\vec{p}^j)) \right).$$

Decomposing $\lambda(\mathcal{P}(\vec{p}^k))$ as well, we find the expression

$$\lambda(\mathcal{P}_{d_{k-1}}) \otimes \circ_{p_1}^{\vee, k-1} \otimes \cdots \otimes \circ_{p_{d_{k-1}}}^{\vee, k-1} \otimes \lambda(\mathcal{P}_{d_k}) \otimes \circ_{p_1}^{\vee, k} \otimes \cdots \otimes \circ_{p_{d_k}}^{\vee, k} \otimes \lambda(Q)$$

in the right-hand side. Using the identification of \mathcal{P}_d with \mathcal{R}_{d+1} discussed in Remark (4.2), and Lemma 8.3, we may replace this by

$$\lambda(\mathbb{R}) \otimes \lambda(\mathcal{P}(\vec{p}^{k-1} \cup \vec{p}^k))$$

at the cost of a sign given by

$$(8.39) \quad d_{k-1} + (d_k + 1) \deg(\vec{p}^{k-1}).$$

Note that in Lemma 8.3, d_1 plays the role of $d_{k-1} + 1$, d_2 that of $d_k + 1$, and k is our d_{k-1} (making for incompatible notations). We now move the normal co-vector $\lambda(\mathbb{R})$ to the beginning of the right-hand side, using Lemma 8.4 to cancel with the normal vector of the inclusion of \mathcal{S}_{r-1} into \mathcal{S}_r in which the k th and $(k-1)$ st incoming leaves are collapsed. The resulting sign is

$$(8.40) \quad k + r + \deg(x_0) + \sum_{j=1}^{k-2} (1 + d_j + \deg(\vec{p}^j) + n)$$

The final output of this procedure is an isomorphism

$$(8.41) \quad \lambda(\mathcal{C}(x_0, \vec{p})) \otimes \lambda(Q^{r-1}) \cong \\ \lambda(\mathcal{S}_{r-1}) \otimes \lambda(W^s(x_0)) \otimes \lambda(\mathcal{P}(\vec{p}^1)) \otimes \cdots \otimes \lambda(\mathcal{P}(\vec{p}^{k-1} \cup \vec{p}^k)) \otimes \cdots \otimes \lambda(\mathcal{P}(\vec{p}^r)),$$

which is exactly the analogue of the isomorphism of Equation (8.37) for case (4). In particular, the difference between the orientations of a

singular shrub as a boundary in cases (4) and (3) is given by the sum of the expressions appearing in Equations (8.38), (8.39), and (8.40):

$$(8.42) \quad 1 + d_k \deg(\bar{p}^{k-1}) \\ + (n+1) \left(k + r + \deg(x_0) + \sum_{j=1}^{k-1} (1 + d_j + \deg(\bar{p}^j)) \right).$$

We now consider the effect of the sign changes encoded in Equation (8.36). First, we note that the difference between $\dagger(\bar{\mathcal{P}})$ on either side is given by

$$(8.43) \quad (k-1) + \sum_{j=k}^r \deg(\mathcal{P}(\bar{p}^j)),$$

from which we conclude that the second term changes by

$$(8.44) \quad (n+1) \left(k + \sum_{j=k}^r (1 + d_j + \deg(\bar{p}^j)) \right).$$

Since the moduli space we are studying is 1-dimensional, we have an equality (modulo 2)

$$(8.45) \quad \deg(x_0) = d + \sum_{j=1}^r \deg(\bar{p}^j),$$

which readily implies that (8.44) cancels the third term of Equation (8.42). The reader may easily check that the remaining terms of Equation (8.36) exactly cancel with the first two terms of Equation (8.42), completing our argument.

Appendix A. The case of clean intersections

In this section, we shall show that our construction, which so far yields Theorem 1.3 whenever B is a point, can be extended to the general case.

The starting point is the observation that, whenever two Lagrangian manifolds Q_1 and Q_2 meet cleanly along a submanifold B , we have a perfect pairing on the normal bundles $N_{Q_1}B$ and $N_{Q_2}B$; in particular, these normal bundles are dual, and hence isomorphic as unoriented vector bundles. Conversely, given smooth embeddings of a compact manifold B into Q_1 and Q_2 with normal bundles isomorphic to a given vector bundle NB , the choice of such an isomorphism determines a plumbing of T^*Q_1 and T^*Q_2 along a neighborhood of B .

Explicitly, we fix a Riemannian metric on B and on the vector bundle NB , which induces a metric on the total space. We may assume that we are given a Riemannian metric on Q_i whose restriction to a neighborhood of B is isometric to the disc bundle of radius 4 in NB , with U_i

corresponding to the disc bundle of radius 2. By Weinstein's theorem, there is a neighborhood of B

$$(A.1) \quad M_B \subset T^*Q_i$$

which is symplectomorphic to the disc bundle of a symplectic vector bundle over a neighborhood of the zero section in T^*B ; this vector bundle is isomorphic to the pull-back of $NB \otimes \mathbb{C}$ to the cotangent bundle of B . By construction, $NB \otimes \mathbb{C}$ carries two Lagrangian sub-bundles obtained as the inclusion of the real part NB and the imaginary part $\sqrt{-1}NB$. The Weinstein chart may be chosen so that the image of Q_i is the union (over all points in B) of the real parts. We construct a manifold

$$(A.2) \quad M = D^*Q_1 \#_{D^*B} D^*Q_2$$

by gluing disc cotangent bundles of Q_1 and Q_2 along this common subset after complex multiplication in $NB \otimes \mathbb{C}$. As this map is an exact symplectomorphism for the appropriate Liouville form, M is again an exact symplectic manifold with exact Lagrangian embeddings of Q_1 and Q_2 intersecting cleanly along B . In fact, we can choose a Liouville form on M which away from B agrees with the canonical forms on D^*Q_i and which near B can be written as a sum

$$(A.3) \quad \theta_{T^*B} \oplus \theta_{NB},$$

where the first form the pull-back of the canonical form on the cotangent bundle and the second defines a flow which radially rescales each fiber.

A.1. Fukaya category. The fiberwise squared norm of a point in M_B (considered as a vector in either T^*Q_1 or T^*Q_2) defines a map

$$\rho: M_B \rightarrow [0, +\infty)^2.$$

As in Equation (2.41), we define

$$M_B^{\text{in}} = (\chi \circ \rho)^{-1}(\epsilon).$$

This is a smooth manifold whose boundary is transverse to the Liouville flow of (A.3), and which away from a neighborhood of B agrees with the ϵ -neighborhood of the zero section in T^*Q_i . As before, we write

$$M^{\text{in}} \subset M$$

for the Liouville subdomain obtained by extending this submanifold to all of M . Condition (2.43) now makes sense in the clean intersection case as well, and the construction of the Fukaya category proceeds as in the transverse case.

The analogue of Condition (2.42) is the requirement that the Floer complexes $CF^*(Q_1, Q_2)$ vanish in degrees smaller than 0 or larger than $\dim(B)$. This is easy to achieve if B is connected by shifting the grading on Q_1 . Otherwise, one needs to pick an appropriate complex volume form on the plumbing in order for this result to hold (using the

fact that the space of complex volume forms is an affine space over $H^1(M, \mathbb{Z})$, and that $H^0(B, \mathbb{Z})$ controls the difference between this space and $H^1(Q_1, \mathbb{Z}) \oplus H^1(Q_2, \mathbb{Z})$.

However, if we were considering the situation of Theorem 1.3, the total space W may not admit any complex volume form for which the Floer cohomology groups are supported in such a degree. In this case, there are integers $\{m_k\}_{k=1}^r$ labeled by the components of B such that the Floer cohomology $HF^*(Q_1, Q_2)$ splits as a direct sum, with the summand corresponding to B^k supported in degrees lying in the interval $[m_k, m_k + \dim(B^k)]$. With appropriate choices of perturbation data, we may ensure that the Floer complexes also admit such a splitting.

A.2. Morse category. Let $Q_{1,2}$ and $Q_{2,1}$, respectively, denote the closures of U_1 and U_2 in Q_1 and Q_2 ; these are diffeomorphic to the disc bundle associated to NB , and the plumbing construction fixes such a diffeomorphism. Note that the construction of the Morse category in Section 2.2 did not make any assumption about U_i (other than the fact that their closures are smooth manifolds with boundary). In particular, the construction may be extended to the case of clean intersections.

However, whenever $Q_{1,2}$ is disconnected, we need a generalization which accounts for the fact Floer cochains are not necessarily supported in degree $[0, k]$ and $[n - k, n]$ for a fixed integer k . We choose integers $\{m_k\}_{k=1}^r$ indexed by the component of $Q_{1,2}$ and define morphism spaces

$$(A.4) \quad Hom_*^M(Q_1, Q_2) = \bigoplus_{k=1}^r CM^*(f_{1,2}^k)[m_k],$$

$$(A.5) \quad Hom_*^M(Q_2, Q_1) = \bigoplus_{k=1}^r CM^*(f_{2,1}^k)[-m_k]$$

where $f_{i,j}^k$ is the restriction of $f_{i,j}$ to the k th component of $Q_{i,j}$ and we are using the standard notation from homological algebra for a shift in degree; the grading on $Hom_*^M(Q_i, Q_i)$ is unchanged. The shifts in gradings do not affect the degree of the contribution of each gradient tree to the higher products, so that these groups are morphism spaces in a \mathbb{Z} -graded A_∞ category. To see this, assume for simplicity that we are considering a gradient tree with $i_0 = i_d$. An argument analogous to the one given in the proof of Lemma 2.12 implies that, in this case, every labeled gradient tree that contributes to the Morse category has the property that the numbers of external edges labeled by $f_{1,2}^k$ and $f_{2,1}^k$ agree for every $1 \leq k \leq r$, so that the formula for expected dimension still holds. The generalization of this argument to the case $i_0 \neq i_d$ is left to the reader.

A.3. Simplicial category. Let us now choose simplicial triangulations \mathcal{Q}_i of Q_i and assume that there is a pair of codimension 0 submanifolds

with boundary $N_i \subset Q_i$ which are the closures of top-dimensional simplices, each properly including U_i as a deformation retract. We assume that the diffeomorphism $U_1 \cong U_2$ determined by the plumbing construction extends to a cellular homeomorphism

$$N_1 \cong N_2.$$

Writing \mathcal{N} for the induced subdivision on an abstract manifold N , and $\{\mathcal{N}^k\}_{k=1}^r$ for the triangulations on the components of this space, we define a category $\text{Simp}_B(Q_1, Q_2)$ with objects Q_1 and Q_2 , and morphism spaces

$$\begin{aligned} \text{Hom}_*^S(Q_i, Q_i) &= C^*(Q_i), \\ \text{Hom}_*^S(Q_1, Q_2) &= \bigoplus_{k=1}^r C^*(\mathcal{N}^k)[m_k], \\ \text{Hom}_*^S(Q_2, Q_1) &= \bigoplus_{k=1}^r C^*(\mathcal{N}^k, \partial\mathcal{N}^k)[-m_k]. \end{aligned}$$

The construction of composition maps and the proofs that the appropriate diagrams commute follows word by word the arguments given in the case N_i is the closure of a simplex. We simply observe that the composition maps still have degree 0 because the different components of \mathcal{N} do not interact, and we have chosen changes in gradings which differ by a sign for the morphisms going in opposite directions.

A.4. From simplicial to Morse. The construction of a functor from the simplicial to the Morse category relied on an appropriate choice of dual subdivision. In the clean intersection case, we choose \check{Q}_i and write \check{N}_i for those cells dual to interior cells of \mathcal{N} , replacing Condition (3.2) by the assumption that the cells of \check{Q}_i intersect the boundary of \bar{U}_i transversely, and that we have nested weak homotopy equivalences

$$(A.6) \quad \check{N}_i \subset \bar{U}_i \subset N_i.$$

Whenever B is disconnected, we assume that the choices of shifts in grading in the simplicial and Morse categories agree. All arguments extend with only minor modifications; e.g., in Lemma 3.8, we should require the appropriate neighborhood of the boundary of $Q_{2,1}$ not to intersect any cell of $Q_{2,1}$ which is dual to a cell whose interior lies entirely in N .

A.5. From Floer to Morse. In order to construct a functor from the Fukaya category to the Morse category, we start by choosing the integers m_k in Equation (A.4), whenever B is disconnected, to agree with those coming from Floer theory as discussed in Addendum A.1.

The next step is to construct a foliation of M_B which we shall extend to M . Near each point $b \in B$, M_B splits as the product of the disc

cotangent bundle of B with a domain in $\mathbb{C}^{n-\dim(B)}$. We equip this region with the product of the cotangent foliation of T^*B with the model foliation \mathcal{L} on $D^*L_1\#D^*L_2$ (with \mathbb{C}^n replaced by $\mathbb{C}^{n-\dim(B)}$) constructed in Section 5.

We claim that this foliation of $\mathbb{C}^{n-\dim(B)}$ can be chosen to be invariant under the action of the orthogonal group $O(n - \dim(B))$ embedded in unitary group (in Figure 7, this corresponds to the invariance under rotation by π), which is the structure group for the normal bundle of T^*B in M . To see this, recall that we first constructed a foliation by affine Lagrangian planes in Equation (5.8), which we then deformed, using a Hamiltonian flow. Since only the norm of a point in $\mathbb{R}^{n-\dim(B)}$ enters in the construction in Equation (5.8), the corresponding foliation is tautologically invariant under $O(n - \dim(B))$; one can then easily check that the construction of the deforming Hamiltonian flow can be made invariant, say, by averaging it over the action of the orthogonal group.

We conclude that the foliations defined locally on neighborhoods of points of B are compatible with each other, and define a foliation on a neighborhood of B in M which agrees with the cotangent bundle foliations of T^*Q_i away from a neighborhood of B , and hence extends to all of M . To construct a functor in this case, note that Condition (5.11) makes sense as stated, but that we cannot choose $f_{i,j}$ to be given by the squared length of a normal vector, as this function is only Morse-Bott. However, choosing $f_{i,j}$ to be a C^1 -small generic perturbation of this fiberwise quadratic function, one concludes that, away from a small neighborhood of B , the distance to B increases (or decreases) along the appropriate gradient flow, so that the analogue of Lemma 5.3 holds.

With the choice of Hamiltonian perturbations coming from (5.11), we obtain an identification between generators of Floer and Morse complexes, and we may carry through the remainder of the construction of an A_∞ functors. The discussion of which moduli spaces are counted, as well as issues of transversality, require no modification. As to compactness, note that

$$L_{B,-\Delta} = \cup_{b \in B} (1 - \sqrt{-1})N_b B \subset \cup_{b \in B} N_b B \otimes \mathbb{C}$$

defines a Lagrangian submanifold of the total space of $NB \otimes \mathbb{C}$, and hence a Lagrangian in M locally near B . In analogy with Equation (6.4), we consider the submanifold of M

$$D_{1/2}Q_1 \cup D_{1/2}Q_2 \cup \left(D_{\sqrt{2}}B \cap D_{\sqrt{4-\delta}}L_{B,-\Delta} \right),$$

where the subscripts stand for the radii of various neighborhoods. Working in charts near each point $b \in B$, the discussion of Section 6 shows that the boundary of this manifold with corners is transverse to the Liouville flow, and hence can be smoothed to a Liouville subdomain

$M^{\text{mi}} \subset M$. As long as the parameter ϵ in the construction of the foliation is chosen much smaller than δ , Lemma 5.12, proved in Section 6 by using the contact type condition near ∂M^{mi} and ∂M^{in} , applies in the clean intersection case as well.

References

- [1] M. Abouzaid, *Framed bordism and Lagrangian embeddings of exotic spheres*, available at arXiv:0812.4781.
- [2] ———, *Morse Homology, Tropical Geometry, And Homological Mirror Symmetry For Toric Varieties*, *Selecta Math. (N.S.)* **15** (2009), no. 2, 189–270.
- [3] M. Abouzaid & P. Seidel, *An open string analogue of Viterbo functoriality*, *Geom. Topol.* **14** (2010), no. 2, 627–718, DOI 10.2140/gt.2010.14.627, MR 2602848.
- [4] M. Akaho & D. Joyce, *Immersed Lagrangian Floer theory*, available at arXiv.org:0803.0717.
- [5] J.M. Boardman & R.M. Vogt, *Homotopy invariant algebraic structures on topological spaces*, *Lecture Notes in Mathematics*, **347**, Springer-Verlag, Berlin, 1973, MR 0420609.
- [6] A. Floer, *Monopoles on asymptotically flat manifolds*, The Floer memorial volume, *Progr. Math.*, **133**, Birkhäuser, Basel, 1995, pp. 3–41, MR 1362821.
- [7] S. Forsey, *Quotients of the multiplihedron as categorified associahedra*, *Homology, Homotopy Appl.* **10** (2008), no. 2, 227–256.
- [8] K. Fukaya & Y.-G. Oh, *Zero-loop open strings in the cotangent bundle and Morse homotopy*, *Asian J. Math.*, **1**, (1997), no. 1, 96–180, MR 1480992.
- [9] K. Fukaya, Y.-G. Oh, H. Ohta, & K. Ono, *Lagrangian intersection Floer theory: Anomaly and obstruction*, *AMS/IP Studies in Advanced Mathematics*, **46**, American Mathematical Society, Providence, RI, 2009, MR 2310372.
- [10] A. Hatcher, *Algebraic topology*, Cambridge University Press, Cambridge, 2002, MR 1867354.
- [11] M. Kontsevich & Y. Soibelman, *Homological mirror symmetry and torus fibrations*, *Symplectic geometry and mirror symmetry* (Seoul, 2000), World Sci. Publ., River Edge, NJ, 2001, pp. 203–263, MR 1882331.
- [12] P. Kronheimer & T. Mrowka, *Monopoles and three-manifolds*, *New Mathematical Monographs*, **10**, Cambridge University Press, Cambridge, 2007, MR 2388043.
- [13] D. McDuff & D. Salamon, *J-holomorphic curves and symplectic topology*, *American Mathematical Society Colloquium Publications*, **52**, American Mathematical Society, Providence, RI, 2004, MR 2045629.
- [14] J. Milnor, *Differential topology*, *Lectures on Modern Mathematics*, **II**, Wiley, New York, 1964, pp. 165–183, MR 0178474.
- [15] M. Poźniak, *Floer homology, Novikov rings and clean intersections*, *Northern California Symplectic Geometry Seminar*, *Amer. Math. Soc. Transl. Ser. 2*, **196**, Amer. Math. Soc., Providence, RI, 1999, pp. 119–181, MR 1736217.
- [16] W.-D. Ruan, *The Fukaya category of symplectic neighborhood of a non-Hausdorff manifold*, available at math.SG/0602119.

- [17] P. Seidel, *Fukaya categories and Picard-Lefschetz theory*, Zurich Lectures in Advanced Mathematics, European Mathematical Society (EMS), Zürich, 2008, MR 2441780.
- [18] ———, *A long exact sequence for symplectic Floer cohomology*, *Topology*, **42**, (2003), no. 5, 1003–1063, MR 1978046.
- [19] J. D. Stasheff, *Homotopy associativity of H-spaces*, I, II, *Trans. Amer. Math. Soc.* **108**, (1963), 275–292; *ibid.* **108** (1963), 293–312, MR 0158400.

MIT DEPARTMENT OF MATHEMATICS
77 MASSACHUSETTS AVE
CAMBRIDGE, MA 02193 USA
E-mail address: abouzaid@math.mit.edu