Enchiridion: Mathematical User's Guides
`mathusersguides.com`
Volume 3 (September 2017)

A user's guide: A monoidal model for Goodwillie derivatives

Sarah Yeakel

Contents

1. Key insights and central organizing principles

1.1. The question. In order to understand certain functors, Goodwillie developed a theory of polynomial approximations for homotopy invariant functors from pointed topological spaces $\mathcal{T}$ to the categories of spaces $\mathcal{T}$ or spectra $\mathcal{S}\mathrm{p}$ in a series of papers called Calculus I, II, and III [**Goo90**, **Goo91**, **Goo03**], a nod to the Taylor series approximation method of function calculus.

He describes a way to canonically assign to a functor F a sequence of "polynomial" (called *n-excisive*) functors P_nF approximating F, which fit into a tower of fibrations analogous to a Taylor series. Thus this tower of functors along with the natural maps $p_nF : F \to P_nF$ are called the Taylor tower for F:

$$\begin{tikzcd}
F(X) \arrow[dr, bend left] \arrow[drr, bend left] & & & & & \\
\cdots \arrow[r] & P_nF(X) \arrow[r] & P_{n-1}F(X) \arrow[r] & \cdots \arrow[r] & P_1F(X) \arrow[r] & P_0F(X) \simeq F(*)
\end{tikzcd}$$

As in function calculus, one tries to study the functor F by studying its Taylor tower, and this is a good approximation when F is *analytic*, which implies $F(X) \simeq \operatorname{holim}_n P_nF(X)$ for sufficiently connected X; we call this convergence in a radius of convergence. Many functors are analytic; some examples include the identity functor of spaces, the functor $F(X) = \Sigma^\infty X^{\wedge k}$, and the representable functors $F(X) = \operatorname{Hom}(K, X)$ for finite dimensional spaces K.

While the analogy with function calculus is fun and illuminating, there is an immediate roadblock to understanding functor calculus with this approach because the n-excisive approximations, P_nF, are difficult to compute; for example, the first approximation of the identity functor is $P_1Id(X) \simeq \Omega^\infty\Sigma^\infty X$, the stable homotopy functor. Since the Taylor tower of the identity functor converges (on appropriately connected spaces), we can view the levels P_nF as interpolating between stable (at level 1) and unstable (in the limit) homotopy theory, which is all rather nontrivial to compute for many spaces.

Because the levels of the Taylor tower are computationally opaque, we turn to the homotopy fibers $D_nF = \operatorname{fiber}(P_nF \to P_{n-1}F)$, or *layers*, of the Taylor tower, with the hopes that the polynomial parts can be reconstructed once the layers are known.

One can think of the fiber D_nF as a difference of the nth polynomial approximation from the $n-1$st; in function calculus, this procedure would produce the n-homogeneous functions $D_nf = \frac{f^{(n)}(0)\cdot x^n}{n!}$. Indeed, the analogy is again justified by Goodwillie's classification of the layers.

Theorem 1.1 ([**Goo03**]).

$$D_nF(X) \simeq \Omega^\infty \left(\partial_nF \wedge X^{\wedge n}\right)_{h\Sigma_n}$$

where ∂_nF is a spectrum with Σ_n-action called the n-th derivative *of F.*

The symmetric group Σ_n acts on the smash product by permuting the factors and $(-)_{h\Sigma_n}$ denotes the homotopy orbits, which could be interpreted as dividing by $n!$. Taken together, the derivatives form a symmetric sequence in the category of spectra, and we see that the layers of the tower are determined by ∂_*F. This means that perhaps the extension problems of recovering the levels of the Taylor tower from the layers could be solved by exploring extra structure in the derivatives. Indeed, the derivatives are rich in structure; it should be noted that many people have worked on this problem, to much avail. As a start, Goodwillie identifies the homotopy type.

Theorem 1.2 ([**Goo03**]). *The n-th derivative of F is equivalent to the multilinearization of the nth cross-effect.*

$$(\Omega^\infty)\partial_nF \simeq \operatorname*{hocolim}_{k_1,\ldots,k_n\to\infty} \Omega^{k_1}\cdots\Omega^{k_n} cr_nF(\Sigma^{k_1}S^0, \ldots, \Sigma^{k_n}S^0)$$

The Σ_n-action is induced by permuting the variables of cr_nF; in the multilinearization, this also permutes the loops. The nth cross effect is a functor of

n variables which can be thought of as a measurement of the failure of F to be degree $n-1$ (in a sense). For example, $cr_1F(X) = \mathrm{hofib}(F(X) \to F(*))$, so if F is degree 0 (or constant), cr_1F is trivial. The higher cross effects are defined as total homotopy fibers of cubes, for example, cr_2F is a functor of two variables, defined as the total fiber of the following 2-cube.

$$\begin{tikzcd}
F(X \vee Y) \arrow[r] \arrow[d] & F(X) \arrow[d] \\
F(Y) \arrow[r] & F(*)
\end{tikzcd}$$

Considering all the derivatives of a functor at the same time yields a symmetric sequence in spectra. Thus we may think of the derivatives as a functor

$$\partial_* : \mathrm{Fun}(\mathcal{T}, \mathcal{T}) \to \mathrm{Fun}(\Sigma, \mathcal{S}\mathrm{p}).$$

Both the domain and codomain have monoidal structures; in the domain, it's composition of functors, in the codomain, the circle product of symmetric sequences. This point of view leads to a question posed by Arone and Ching in the introduction of [**AC11**], which is the main goal of [**Yea17**].

Question 1.1. *Is ∂_* monoidal?*

Since monoidal functors preserve monoids, this would imply that the derivatives of monads (which are monoids in the category of endofunctors) on topological spaces are operads (which are the monoids under the circle product of symmetric sequences). Then the derivatives of the identity functor would form an operad, and derivatives of other functors would be modules over that operad. This is the kind of extra structure we are looking for, although Arone and Ching have shown [**AC11**] that it is not enough to recover the levels of the Taylor tower. In nontechnical terms, being monoidal says that the derivatives of functors behave well with respect to composition and the derivatives of the identity functor fit together with themselves and with other derivatives in interesting ways.

Specifically, being monoidal would require a natural transformation $\partial_*F \circ \partial_*G \to \partial_*(F \circ G)$ and a map $\mathbb{S} \to \partial_1 Id$, where the first $\circ$ is the circle product of symmetric sequences and the second is composition of functors. These maps should fit into commutative diagrams for commutativity, unitality, and equivariance. The first map looks like a sort of chain rule for functors, describing a relationship between the derivative of a composition and the derivatives of the composites.

In [**Yea17**], we describe a model for the derivatives which is monoidal, giving a positive answer to the question of Arone and Ching above. Verifying that all the necessary diagrams commute is part of the drama, and we will now relive some of the obstacles encountered, drawing attention to some important concepts.

1.2. The hurdles. There are some drawbacks to working with Goodwillie's model for the derivatives. By universal properties and choosing models for homotopy (co)limits, it seems that a map on compositions $\mu : \partial_* F \circ \partial_* G \to \partial_*(F \circ G)$ is not hard to construct, but such a map would not necessarily be strictly associative. That is, one could easily define a map μ which is associative and unital up to homotopy, but strict associativity requires a different model for the derivatives. While frustrating, this leads to a very important theme of the paper.

KEY IDEA 1.2. *The nitpicky details are important; be careful choosing models for homotopy functors.*

A map of symmetric sequences is a levelwise morphism for each natural number. In level n, μ takes the form

$$\bigvee_{j_1+\cdots+j_k=n} \partial_k F \wedge \partial_{j_1} G \wedge \cdots \wedge \partial_{j_k} G \longrightarrow \partial_n(F \circ G).$$

A map out of a coproduct can be defined by a product of maps out of the summands, and some of these summands, for example when $n = 2$, look like $\partial_1 F \wedge \partial_2 G$, $\partial_2 F \wedge \partial_1 G \wedge \partial_1 G$, or $\partial_2 F \wedge \partial_0 G \wedge \partial_2 G$ (the coproduct is enormous). In the paper, we focus on *reduced* functors, those which preserve the one point space, and thus the 0th derivative is the value of the functor at a point, which is trivial. This makes the summands with 0th derivatives uninteresting. Thus on level one, μ boils down to a map $\partial_1 F \wedge \partial_1 G \to \partial_1(F \circ G)$.

We will focus on the impossibilities in level one, on the first derivatives. Recall that the first derivative is the linearization of the first cross effect, and linearization is a homotopy colimit

$$\operatorname{hocolim}\left(X \to \Omega\Sigma X \to \Omega^2\Sigma^2 X \to \Omega^3\Sigma^3 X \to \cdots\right)$$

and so the map we desire must combine two of these homotopy colimits into one:

$$\operatorname*{hocolim}_k \Omega^k cr_1 F(S^k) \wedge \operatorname*{hocolim}_\ell \Omega^\ell cr_1 G(S^\ell) \to \operatorname*{hocolim}_n \Omega^n cr_1(F \circ G)(S^n)$$

This is a classical problem in homotopy theory, exactly the roadblock to finding a good category of spectra with a strictly associative smash product and also the cause of delay in defining THH. One solution to this problem, introduced in one form by Bökstedt in the 80s and in another by Smith in the 90s, is to include more symmetry, a technique we can employ by changing the shape of the homotopy colimit. Because of the recurring nature of this issue and resolution, we deem this the major key idea of the project.

KEY IDEA 1.3. *Including symmetry into a straight line homotopy colimit allows for a strictly associative combination.*

The current shape is indexed by the category $\mathbb{N}$ of finite sets $\underline{n} = \{1, 2, \ldots, n\}$ with standard inclusions $\{1, \ldots, n\} \hookrightarrow \{1, \ldots, n, n+1\}$

$$\begin{tikzcd} \emptyset \arrow[r] & \underline{1} \arrow[r] & \underline{2} \arrow[r] & \underline{3} \arrow[r] & \cdots \end{tikzcd}$$

A new shape we could use is indexed by the category $\mathbb{I}$ of finite sets $\underline{n}$ with all injective maps

$$\begin{tikzcd}
\emptyset \arrow[r] & \underline{1} \arrow[r] & \underline{2} \arrow[r] \arrow[loop above, "\Sigma_2"] & \underline{3} \arrow[r] \arrow[loop above, "\Sigma_3"] & \cdots
\end{tikzcd}$$

This category has a symmetric monoidal product given by disjoint union of sets which yields a map on homotopy colimits

$$\operatorname*{hocolim}_{(U,V)\in\mathbb{I}\times\mathbb{I}} G(U,V) \to \operatorname*{hocolim}_{U\amalg V\in\mathbb{I}} G(U \amalg V).$$

A similar map exists for the category $\mathbb{N}$, given by addition, but this does not translate to a strictly associative map on loops (which is why Moore loops were invented).

Indeed, linearization fits into the $\mathbb{I}$ diagram shape because $\Omega^n\Sigma^n$ has a natural Σ_n-action given by permuting the sphere coordinates. Thus, two homotopy colimits over $\mathbb{I}$ can be (strictly!) associatively reindexed to a single homotopy colimit over $\mathbb{I}$.

The next hurdle is that the cross effects must be combined. That is, the cross effects functor

$$cr_* : \mathrm{Fun}(\mathcal{T},\mathcal{T}) \to \mathrm{Fun}(\Sigma, \mathrm{Fun}(\mathcal{T}^*,\mathcal{T}))$$

needs to be monoidal also. Again, this requires strict associativity, and cross effects are defined as total homotopy fibers, so the key to proving this is by careful choice of models, a topic which receives significant airtime in [**Yea17**]. We end up finding that induction works here, so another key idea is the following.

Key Idea 1.4. *A total homotopy fiber of a cube can be computed as an iterated homotopy fiber.*

This is well-known [**Goo91**], but in the case of the cross effects, we must rewrite a total homotopy fiber (an nth cross effect) as an iteration of very particular homotopy fibers: the first cross effects of some functor.

Combining the cross effects in this monoidal way also requires them to have assembly maps in each variable, a detail which is crucial to the models. A functor F has *assembly* if there are maps $F(X) \wedge Y \to F(X \wedge Y)$ for all spaces X, Y. We restrict our attention to *continuous* functors, which have assembly maps that also carry through to the cross effects. Continuity is the condition that there is a continuous map of spaces $\mathrm{Hom}(X,Y) \to \mathrm{Hom}(F(X),F(Y))$, and so the functor F must be pointed (that is, $F(*) = *$). In the category of simplicial sets, a reduced functor ($F(*) \simeq *$) can be replaced functorially with a pointed functor.

The final hurdle we will discuss arises when looking at the higher level maps, for example, $\mu : \partial_1 F \wedge \partial_2 G \to \partial_2(F \circ G)$. The map μ must be equivariant with respect to the Σ_2-actions while still being associative. We will talk about the solution of this problem in Topic 2, but suffice it for now to say

Key Idea 1.5. *The sphere operad of* [**AK14**] *is awesome.*

Essentially, in order for the homotopy colimit combination to be effective for the single hocolim of $\partial_1 F$ with the double hocolim of $\partial_2 G$, we need to double up the sphere coordinates that we are linearizing along in $\partial_1 F$. This requires maps like $\Omega^k S^k \to \Omega^{2k} S^{2k}$, which can be easily defined as suspending by a k-sphere. For the equivariance, we need to choose a perpendicular complement of the existing k-spheres S^k and suspend in that direction. For the associativity to still hold, we need a really good choice of sphere complement, and the sphere operad is exactly that. We'll discuss how to picture this in Topic 2.

2. Metaphors and imagery

We will use this section to give some intuition for homotopy colimits and what "combining" them entails. We will discuss why we should believe that the derivatives that show up in functor calculus are multilinearizations of cross effects, and, as foretold, we will describe the sphere operad and the salient properties that allow for monoidal derivatives.

2.1. Homotopy colimits. The colimit of a diagram is an object that encodes the limiting behavior of maps, and one can think of the homotopy colimit as physically keeping track of the maps along the way. Bousfield and Kan give a description of the homotopy colimit as the geometric realization of a simplicial set indexed by the nerve of the diagram category; for a diagram $\mathcal{X} : \mathcal{C} \to \mathcal{T}$, the n-simplices of the homotopy colimit $\operatorname{hocolim}_{\mathcal{C}} \mathcal{X}$ are $\coprod_{c_0 \to \cdots \to c_n} \mathcal{X}(c_0)$.

Upon realization, each object of the diagram gets a space, each arrow of the diagram gets a copy of the source space times a 1-cell that traces out the map between the corresponding source and target spaces, each pair of composable morphisms gets the source space times a 2-cell which fills in the triangle formed by the composite 1-cells, etc. Two basic diagram shape examples are given below, shown with the homotopy colimits living above their diagrams.

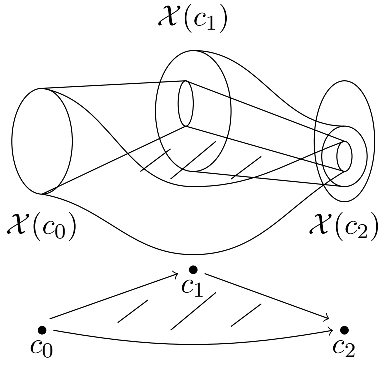

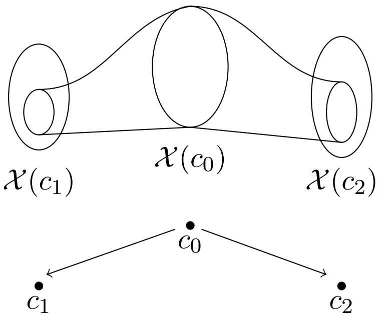

These images are not new to anyone familiar with Dugger's excellent primer [**Dug**], and I strongly recommend it for more information. One can imagine the homotopy colimit space as living above the diagram shape category $\mathcal{C}$, and this gives a bit more intuition when it comes to visualizing maps between homotopy colimits resulting from changing the diagram shape. We will use this imagery to describe the "combining" of homotopy colimits mentioned in Topic 1. Since Bousfield-Kan's description of the homotopy colimit [**BK72**] is the realization of a simplicial set indexed on the nerve of the diagram shape category $\mathcal{C}$, we can think of individual n-cells of that simplicial set as $(x, c_0 \to \cdots \to c_n) \in \mathcal{X}(c_0) \times N_n\mathcal{C}$. In the figure above, for example, a 1-cell is a point in one of the spaces along with its image under an arrow.

An important fact of homotopy colimits is that when there is a map between diagram categories $\alpha : \mathcal{C} \to \mathcal{D}$, there is an induced map between the homotopy colimits indexed over those categories, $\operatorname{hocolim}_{\mathcal{C}} \mathcal{X} \circ \alpha \to \operatorname{hocolim}_{\mathcal{D}} \mathcal{X}$. The simplex $(x; c_0 \to \cdots \to c_n)$ with $x \in \mathcal{X}(\alpha(c_0))$ is sent to the simplex $(x; \alpha(c_0) \to \cdots \to \alpha(c_n))$. A relevant example is the map of diagram categories $\alpha : \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ given by addition. On homotopy colimits, this yields a map $\operatorname{hocolim}_{\mathbb{N}\times\mathbb{N}} \mathcal{X} \circ \alpha \to \operatorname{hocolim}_{\mathbb{N}} \mathcal{X}$. A (truncated) picture is given below, where the first image is a homotopy colimit over $\mathbb{N} \times \mathbb{N}$ and the second is its image under α.

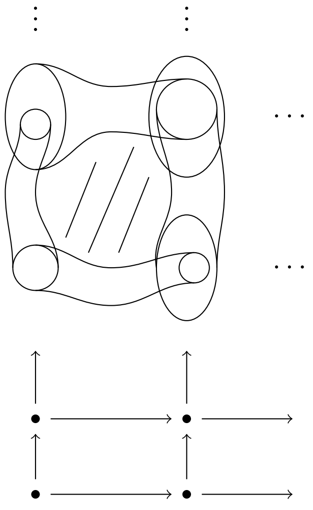

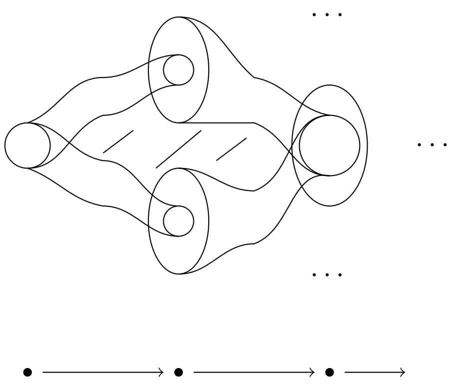

This is what we mean by combining homotopy colimits; the diagram map $\mathbb{N}\times\mathbb{N} \to \mathbb{N}$ given by addition yields a map on homotopy colimits that is associative up to homotopy, but not strictly associative. The problem is similar to that for finding a strictly associative smash product for spectra; there is essentially an $\mathbb{N}\times\mathbb{N}$ grid of spaces (each spectrum has $\mathbb{N}$ spaces, but with structure maps a little different than the grid above) and the goal is to combine this information into only $\mathbb{N}$ spaces. This is possible by choosing a path through the grid, but this is only associative up to homotopy. To make the combination strictly associative, one thing that works for the smash product is to introduce symmetric group actions into the objects of the grid. We are essentially doing the same thing by including all injections of finite sets into the diagram shape, changing the diagram from $\mathbb{N}$ to $\mathbb{I}$. But how does including symmetry help? For the smash product, including the symmetry is like including all possible paths; for the colimits, the extra maps encoding the symmetry make the homotopy colimits enormous, so that they contain all possible ways you could squish the simplices over to one copy of $\mathbb{I}$.

2.2. Function derivatives to Goodwillie derivatives.

While the Goodwillie-Taylor tower and the layers have natural analogs in the function calculus world, the literature does not have a description of why the multilinearized cross effects make any sense as derivatives. This yields a fun exercise in the analogies of calculus, so we include it here.

We will start with what we know from function calculus, the first derivative of a function f at 0,

$$f'(0) = \lim_{x\to 0} \frac{f(x) - f(0)}{x - 0}.$$

We use the standard dictionary to get to functor calculus, so replace x with space X, 0 with the 0-space $*$, and subtraction with homotopy fiber. We also note that the numerator is a familiar object, the first cross-effect.

$$f'(*) = \lim_{X\to *} \frac{\mathrm{hofib}[f(X) \to f(*)]}{\mathrm{hofib}[X \to *]} \simeq \lim_{X\to *} \frac{cr_1 f(X)}{X}$$

To take a limit as a space approaches another space, we need to talk about the "topology" on the category of spaces. The idea is to think of connectivity as giving distance, so being more highly connected means being closer to the one-point space. One way to take a limit as x approaches a is to look at a sequence converging to a, so we choose the convenient sequence of spheres, $S^0, S^1, S^2, \ldots$ which 'converges' to the contractible S^∞. Thus our limit becomes

$$f'(*) \simeq \lim_{n\to\infty} \frac{cr_1 f(S^n)}{S^n} \simeq \lim_{n\to\infty} \Omega^n cr_1 f(S^n).$$

The final equivalence is thinking of looping as the inverse of suspending. It's not a perfect dictionary, but this is the linearization of the first cross effect, as promised. Then the higher derivatives are iterates of this. Recall that the second cross effect is a total homotopy fiber of a square that can be computed with iterative fibers.

$$\begin{tikzcd}
cr_2F(X,Z) \arrow[d, dotted] & & \\
{[F(X \vee Z) - F(X)]} \arrow[r, dashed] \arrow[d, dotted] & F(X \vee Z) \arrow[r] \arrow[d] & F(X) \arrow[d] \\
{[F(Z) - F(*)]} \arrow[r, dashed] & F(Z) \arrow[r] & F(*)
\end{tikzcd}$$

Cross effects existed before functor calculus, and the second cross effect $cr_2 f$ is a function of two variables defined by $cr_2 f(x,z) = [f(x+z) - f(x)] - [f(z) - f(0)]$.

Then the second derivative at zero is the following.

$$f''(0) = \lim_{x \to 0} \frac{f'(x) - f'(0)}{x - 0} = \lim_{x \to 0} \frac{\lim\limits_{y \to x} \frac{f(y) - f(x)}{y - x} - \lim\limits_{z \to 0} \frac{f(z) - f(0)}{z - 0}}{x - 0}$$

Replace y with $x + z$ and replace the limit over y with a limit as z goes to zero. In our translation, addition is replaced by wedge sum, so we have

$$f''(*) \simeq \lim_{X \to *} \lim_{Z \to *} \frac{f(X \vee Z) - f(X) - f(Z) + f(*)}{XZ} \simeq \lim_{n \to \infty} \lim_{k \to \infty} \Omega^n \Omega^k cr_2 f(S^n, S^k)$$

the multilinearized second cross effect.

2.3. The sphere operad.

We mentioned that the sphere operad of [**AK14**] is a vital component of the definition of monoidal derivatives, and we will describe why it is necessary and how one can think of it. Recall that a goal is to define a map $\partial_1 F \wedge \partial_2 G \to \partial_2(F \circ G)$. As defined in the paper, $\partial_1 F \wedge \partial_2 G$ is

$$\operatorname*{hocolim}_{n} \Omega^n cr_1 F(S^n) \wedge \operatorname*{hocolim}_{k,\ell} \Omega^k \Omega^\ell cr_2 G(S^k, S^\ell).$$

The idea is to double up the n-coordinates so that once the second cross effect is assembled into the first, the S^n can assemble into both variables of $cr_2 G$. This requires a map $\Omega^n H(S^n) \to \Omega^n \Omega^n H(S^n \wedge S^n)$ which is associative and equivariant (with respect to swapping the two copies of S^n). It's easier to envision this if H is the identity functor.

Let S^n be the one-point compactification of $\mathbb{R}^n$, and consider an element f of ΩS^1, a map $f : S^1 \to S^1$. To get a map $S^2 \to S^2$, we could use $S^1 \wedge f$, but this new element of $\Omega^2 S^2$ needs to have a trivial Σ_2-action, so must be symmetric on the two sphere coordinates, not just the identity on one. One fix is to put the element f in as the diagonal in $\mathbb{R}^2$ and suspend it in the antidiagonal direction. That is, given a map

$$f : .... \to - - - - -$$

the options are to think of this sphere as one of the coordinates of S^2 while suspending in the direction of the other coordinate, or to think of it as the diagonal in S^2 while suspending in the perpendicular direction (these two options are pictured below).

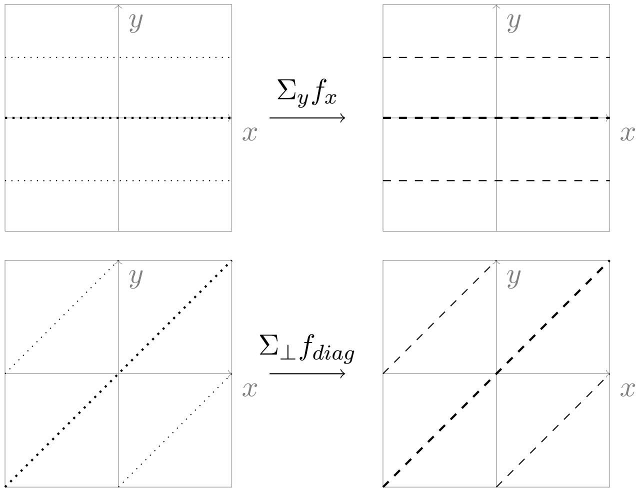

The suspension along the second coordinate, depicted in the top picture, is clearly not equivariant with respect to permuting coordinates; the second suspension, while equivariant in the appropriate way, is not strictly associative. This can be seen by considering the suspension that would define maps $\Omega S^1 \to \Omega^k S^k$ for higher k. This is where the sphere operad comes into play. We are essentially finding a "complement" sphere to suspend along; in the first example, we've chosen one along the vertical axis, and in the second example, the complement runs along the antidiagonal $y = -x$. We will describe the sphere operad and how its components give good complements.

First, we describe the nonunital simplex operad, whose nth space is an open $n-1$-dimensional simplex. In level n, it is the limit over k of the n-simplices in $\mathbb{R}^{n+1}$ described by $x_1^k + \cdots + x_n^k = 1$ with $x_i > 0$. In level 2, this looks something like the final image below.

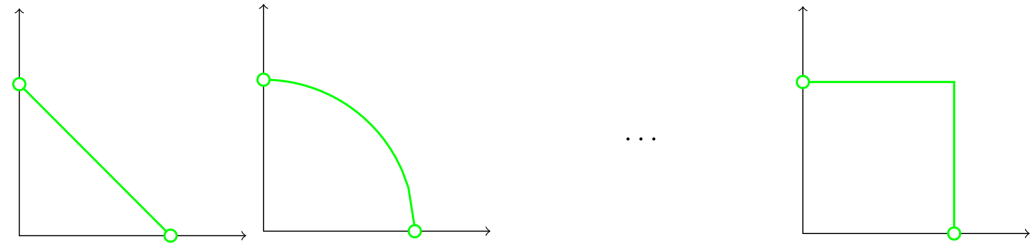

$x_1 + x_2 = 1$ $x_1^2 + x_2^2 = 1$ $x_1^k + x_2^k = 1$ $x_1^\infty + x_2^\infty = 1$

The sphere operad $\mathbf{S}$ is the (levelwise) one-point compactification of the simplex operad, so the nth space of $\mathbf{S}$ is homeomorphic to S^{n-1}. The operad composition maps are homeomorphisms

$$S^{k-1} \wedge S^{j_1-1} \wedge \cdots \wedge S^{j_k-1} \to S^{j_1+\cdots+j_k-1}.$$

There is a map of operads $\mathbf{S} \to \mathrm{Coend}(S^1)$ such that for each $n \geq 1$ the map $\mathbf{S}_n = S^{n-1} \to \Omega S^n$ is adjoint to a homeomorphism $S^{n-1} \wedge S^1 \to S^n$. Since the Σ_n-action on the coendomorphism operad of S^1 permutes the n coordinates of S^n, this defines a Σ_n-equivariant map $S^1 \wedge \mathbf{S}_n \cong S^n$. Finally, there is a map of operads $Com \to \mathbf{S}$ such that the composite $Com \to \mathbf{S} \to \mathrm{Coend}(S^1)$ is levelwise the canonical map adjoint to the diagonal map $S^1 \to S^n$.

These are the properties that ensure the equivariance that we need. In our basic example $\Omega S^1 \to \Omega^2 S^2$, the desired map puts f into $\mathbb{R}^2$ on the diagonal, but now suspends along $\mathbf{S}_2$ instead of the antidiagonal. The fact that $\mathbf{S}_k$ assemble into an operad means that we will also get the associativity that we need.

3. Story of the development

Many people say math research involves heaps of mistakes, but we don't often talk about them. I'm not going to sugarcoat how much failure this project involved. At the time of this writing, it hasn't made it through peer-review, so stay tuned. The original ideas of this paper came from my advisor, Randy McCarthy. He got the idea of using a homotopy colimit over the category $\mathbb{I}$ from Bökstedt's definition of THH, and he wondered if it could solve some problems in Goodwillie calculus. I spent the first few years of graduate school asking him questions about spectra, functor calculus, and homotopy colimits. We were mostly concerned with the polynomial approximations P_nF. While the homogeneous approximations D_nF were classified by Goodwillie, the polynomial ones have more obtuse classifications. There are results for special cases of functors, and Arone and Ching show that the Taylor tower can be reconstructed from certain information about the derivatives [**AC16**], but an independent classification was wanting.

Thus, we set out on the project of classifying polynomial functors. Randy was fairly confident that it would work, and he led me unknowingly to reprove many results from his paper with Brenda Johnson [**JM03**] but now in a topological setting. When the project was nearly done in summer 2014, I gave two talks about it in Europe. The first was Goodwillie's 60th birthday conference in Dubrovnik. I spent the week before the conference traveling around, stressing out, and proving that some of my results only held for highly connected spaces on a bus through the Croatian countryside. A little fazed, I spoke at the Young Topologists Meeting where a mistake was found by some students in the audience. I spent the rest of the summer slowly pushing around the bits we were unsure about (had I just phrased them wrong?) until they were clearly in contradiction. It took a couple months to convince Randy we were wrong, and I learned a lot more about why it didn't work in this time. We were feeding k-spheres into the main theorem which only held on k-connected spaces. This was an enormous, embarrassing dead end, but I scraped together the measly results I had for posterity.

I was dismayed at having LIED to the experts in my field, my thesis broken, when Randy suggested I learn about operads and show the derivatives of the identity formed one. At the time, operads were a thing I had seen defined in talks, and I was fairly certain that I would never understand them. I was demoralized to work on yet another difficult problem that was already solved by some geniuses in an (at the time, for me) incomprehensible way.

The original goal of the project was to find an operad structure on the derivatives of the identity functor of spaces by applying definitions and using a homotopy colimit over $\mathbb{I}$. The first order of business was to define the cross effects in the right way to get an associative map. They should fit together in cool nesting cubical ways, but somehow I couldn't pin them down in levels higher than 1. So we found an inductive definition of cross effects which seemed to work. The next step was to figure out if the multilinearizations fit into an operad structure. Sometime in here we also hammered out a proof of the chain rule, essentially a fun induction argument that would hold if we could get the operad structure right. I had to figure out what exactly equivariance was for operads. It turns out that there are two equivariance diagrams for operads, which can be consolidated into one.

$$\begin{tikzcd}
\mathcal{O}(k) \otimes \mathcal{O}(j_1) \otimes \cdots \otimes \mathcal{O}(j_k) \arrow[r, "\sigma\otimes\tau_1\otimes\cdots\otimes\tau_k"] \arrow[d, "id\otimes\sigma^*"] & \mathcal{O}(k) \otimes \mathcal{O}(j_1) \otimes \cdots \otimes \mathcal{O}(j_k) \arrow[dd, "\gamma"] \\
\mathcal{O}(k) \otimes \mathcal{O}(j_{\sigma^{-1}(1)}) \otimes \cdots \otimes \mathcal{O}(j_{\sigma^{-1}(k)}) \arrow[d, "\gamma"] & \\
\mathcal{O}(j_{\sigma^{-1}(1)} + \cdots + j_{\sigma^{-1}(k)}) \arrow[r, "\sigma_*(\tau_1,\ldots,\tau_k)"] & \mathcal{O}(j_1 + \cdots + j_k)
\end{tikzcd}$$

Here $\sigma(\tau_1, \ldots, \tau_k)$ is the composite $(\tau_1 \oplus \cdots \oplus \tau_k)\sigma_*(j_1, \ldots, j_k)$, where for $\sigma \in \Sigma_k$, $\sigma(j_1, \ldots, j_k) \in \Sigma_j$ permutes blocks of size j_s according to σ, and for $\tau_s \in \Sigma_{j_s}$ $\tau_1 \oplus \ldots \oplus \tau_k$ denotes the image of $(\tau_1, \ldots, \tau_k)$ under the natural inclusion of $\Sigma_{j_1} \times \ldots \times \Sigma_{j_k}$ into Σ_j.

After dragging my feet for a while, I checked that $\partial_* Id$ satisfied the complicated diagram, and my map worked swimmingly. Success! Relief! It was around this time that we realized that the result could be phrased in terms of the derivatives functor ∂_* being monoidal, so the proofs were overhauled to fit this general setup. I can't pinpoint the exact moment of this revelation, but it was sometime between January and July 2015 according to my email inbox, and one lead is an email thread with Cary Malkiewich setting up a meeting to talk about the project. My impression is that a lot of people knew the derivatives should be monoidal; it's essentially spelled out in the introduction of [**AC11**], in a future work section teeming with insight and possible projects. When I went to Stockholm in December 2015 to talk to Greg Arone about the project, he showed me that I had not checked the right diagram, and my map actually failed to satisfy the diagram for nontrivial τ's. He was clearly right, and the problem lay in the way the linearizations fit together, that is, the loop sphere problem discussed in Topic 2.

Greg told me that a normal fix for my mistake would be to use a diagonal map, but he warned me that this would not work in my case, and he offered another possible solution, the sphere operad. He mentioned another small oversight about reduced-ness of the cross effects that left my functors without assembly maps which is kind of crucial, but I decided that this was easily fixed. I left Stockholm determined to amend errors and understand Greg's solutions.

I spent the next month trying to understand what would happen if I used the diagonal map. It worked; I could not see what the problem was. I spent the month after trying to fit a complicated associativity diagram onto one page (it was at least 3-dimensional) and convinced multiple mathematicians in Urbana that everything worked, then defended my thesis; it was time to check with Greg. He responded with a clear counterexample showing the map was not associative and a friendly admonition to try the sphere operad.

The sphere operad really did work! I had already deposited my thesis using the incorrect map, and I spent a month checking this last alteration meticulously. Finally, done! I could leave Urbana with peace of mind that my thesis was correct. I had a new collaboration to work on and decided to worry about the publications from the thesis later. In the back of my mind there was a little voice saying, "Are you sure? Didn't Greg say something else?" That pesky issue about reduced-ness of the cross effects was an actual problem. When I returned to the question in the fall, I realized that switching to the category of simplicial sets should solve the problem and (hopefully!) allow things to generalize. I could have posted the paper to the Arxiv in fall 2016, but I was held back by fear of the mass of imperfections that plagued the project from the start. I continued to worry over it for another semester, and finally posted the preprint after a week of hiking in the mountains of Montana.

I have given many talks on the subject, in various states of feeling fraudulent, sometimes totally confident in the results, and sometimes terrified that someone would see the flaw I was actively trying to mend. Despite my doubts, there was always at least one person stubbornly optimistic the project would work out, and that was Randy, so I owe him for convincing me to keep working; at some point there was a vague threat that if I didn't finish, I would be obligated to name my firstborn child after him.

4. Colloquial summary

When my calculus students ask me what I do, I respond that my area is algebraic topology, and they wrinkle their noses and say, “you don’t just do harder and harder calculus?”

I remind them that one major theme in calculus is that the tangent line is often a very good approximation of a function (near the point of tangency). If any of my students have stumbled in here: that will be on your exam. We can draw the graph of a line $y = mx + b$ fairly easily, and the output of even non-integer inputs are simple to determine. This is a significant motif in calculus. In the first semester, we say, having difficulties with your function? Approximate it with a line! In the second semester, we say, the line was great but maybe we can do better. How about a parabola? We notice this is a little better of an approximation, while not getting too much more difficult to compute, so we try higher-degree polynomials. This leads to infinitely long polynomials, or power series, which can sometimes agree with the original function. The concept of approximating hard things with easier ones is at the heart of a lot of topology.

In my research, I approximate functors. These are a lot like functions, but instead of taking numbers as input, they take objects in a category. In topology, we often study the category of topological spaces. This is a collection of all spaces, along with continuous maps between them; so, for example, there is a circle in this category, and also the map from the circle to itself that wraps around twice. There is a map from the circle to a single point which just squashes it. In my work, I think of each space as having one point that is special. The category of spaces has some cool properties; like numbers, we can add and multiple spaces. Adding two spaces A and B is easy; you just glue together their special points. We denote the sum by $A \vee B$. Below we demonstrate the sum of two circles, where S^1 is the notation for a circle.

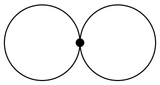

$S^1 \vee S^1$

There are different ways to multiply spaces, but I will only describe one. To multiply space A times space B, at every point in space A, you put a copy of space B. Then you stand back and notice that you could have done it the other way, that is, every point in space B has a copy of space A attached to it. You intuitively know how to do this for basic spaces; it’s likely how you learned multiplication. A two point space $A = \bullet\bullet$ times a three point space $B = \bullet \ \bullet \ \bullet$ is a six point space.

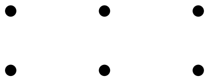

$A \times B$

Now if S^1 is a circle and I is the interval $[0,1]$ on the real line, then $S^1 \times I$ would look like a cylinder.

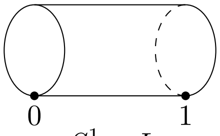

$S^1 \times I$

I encourage you to take a moment before reading further to think about what $I \times I$ looks like. Then try $S^1 \times S^1$ and a three point space times a filled-in circle. Can you think of two spaces that would multiply to a filled-in 3-dimensional cube? What about a ball?

In the category of spaces with a chosen point, there's also an analogy of subtracting spaces, called the fiber. To demonstrate, let's take two circles which have been added, $S^1 \vee S^1$, then subtract one circle. Notice there is actually a map (or function) from $S^1 \vee S^1$ to S^1 in which we squash one circle down to a point and leave the other alone. Let's say we squash the left one. Then the fiber of $S^1 \vee S^1 \to S^1$ is looking at the inverse image of the point. (We're asking, what does the first space have that the second doesn't?) So the fiber is S^1. Similarly, if we took $S^1 \times S^1$ and looked at the map which squashed one copy down, uniformly all the way around the circle, then the fiber at a chosen point is just the inverse image of that point under the map, which is a copy of S^1. So subtraction of spaces is weird; it doesn't act like ordinary numbers. If $Z - X = Y$, then you know that $Z = X + Y$, but here we have two different examples of Z such that fiber $Z - S^1$ is S^1. I encourage you to try to think of more Z's with this property.

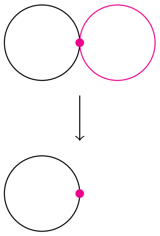

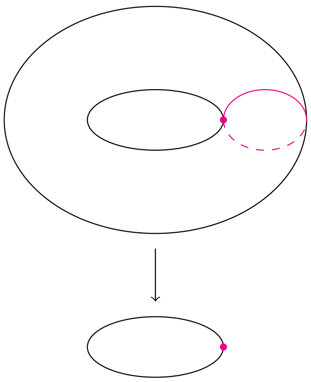

As we've said, functors are a lot like functions, but the domain and range are categories instead of numbers. A function is a way of associating a new number to any number; so $f(x) = x+5$ always adds the number 5 to the input. A functor is like this, but with categories. We can't add 5 to a space, but we can add a space to a space, so we could define a functor from the category of spaces to the category of spaces by $F(X) = X \vee S^1$, which always glues a circle on the chosen point of a space X. There's also a really boring functor, $F(X) = X$, called the identity, and a functor that squashes the entire space to a point, $F(X) = *$, called the trivial functor.

Goodwillie calculus is the approximation of functors with polynomial ones. The definition of polynomial functor is defined in terms of what the functor does to glued together spaces, which means there can be polynomials that don't look like X^n. But a mathematician named Tom Goodwillie defined a degree n polynomial approximation P_nF of a functor F for every n, and he showed that if we subtract $P_{n-1}F$ from P_nF, we always get something like $F' \times X^n$. That's a coefficient space times X multiplied by itself n times, which looks like a standard degree n function. In ordinary calculus, the coefficients on polynomial approximations are given by the higher derivatives of the function, and once you know the derivatives, you know the polynomials because you can just add them up. But because subtraction works differently with spaces, just knowing the derivatives isn't enough to determine the polynomials anymore, so we need to study how the different levels interact. This paper is an effort to understand the derivatives of a functor, their relationship to each other, and how that may influence how the polynomials interact.

References

[AC11] Greg Arone and Michael Ching, *Operads and chain rules for the calculus of functors*, Astérisque **338** (2011), vi+158 (English, with English and French summaries).

[AK14] Greg Arone and Marja Kankaanrinta, *The sphere operad*, Bull. Lond. Math. Soc. **46** (2014), no. 1, 126–132.

[AC16] Greg Arone and Michael Ching, *Cross-effects and the classification of Taylor towers*, Geom. Topol. **20** (2016), no. 3, 1445–1537.

[BK72] A.K. Bousfield and D.M. Kan, *Homotopy Limits, Completions and Localization*, Lecture Notes in Mathematics, vol. 304, Springer-Verlag Berlin Heidelberg, 1972.

[Dug] Dan Dugger, *A primer on homotopy colimits*, http://pages.uoregon.edu/ddugger/hocolim.pdf.

[Goo90] Thomas G. Goodwillie, *Calculus. I. The first derivative of pseudoisotopy theory*, K-Theory **4** (1990), no. 1, 1–27.

[Goo91] ______, *Calculus. II. Analytic functors*, K-Theory **5** (1991/92), no. 4, 295–332.

[Goo03] ______, *Calculus. III. Taylor series*, Geom. Topol. **7** (2003), 645–711 (electronic).

[JM03] B. Johnson and R. McCarthy, *A classification of degree n functors. I*, Cah. Topol. Géom. Différ. Catég. **44** (2003), no. 1, 2–38 (English, with French summary).

[Yea17] S. Yeakel, *A Monoidal Model for Goodwillie Derivatives*, ArXiv e-prints (2017), available at `1706.06915`.

Department of Mathematics, University of Maryland, College Park, MD 20742

E-mail address: `syeakel@math.umd.edu`