

Functor String Diagrams: A Novel Approach to Flexible Diagrams for Applied Category Theory

Vincent Abbott

College of Engineering, Computing and Cybernetics
Australian National University
Canberra, ACT, Australia
Vincent.Abbott@anu.edu.au

Gioele Zardini

Laboratory for Information and Decision Systems
Massachusetts Institute of Technology
Cambridge, MA, USA
gzardini@mit.edu

The study of abstraction and composition—the focus of category theory—naturally leads to sophisticated diagrams which can encode complex algebraic semantics. Consequently, these diagrams facilitate a clearer visual comprehension of diverse theoretical and applied systems. Complex algebraic structures—otherwise represented by a forest of symbols—can be encoded into diagrams with intuitive graphical rules. The prevailing paradigm for diagrammatic category theory are monoidal string diagrams whose specification reflects the axioms of monoidal categories. However, such diagrams struggle in accurately portraying crucial categorical constructs such as functors or natural transformations, obscuring central concepts such as the Yoneda lemma or the simultaneous consideration of hom-functors and products. In this work, we introduce *functor string diagrams*, a systematic approach for the development of categorical diagrams which allows functors, natural transformations, and products to be clearly represented. We validate their practicality in multiple dimensions. We show their usefulness for theoretical manipulations by proving the Yoneda lemma, show that they encompass monoidal string diagrams and hence their helpful properties, and end by showing their exceptional applied utility by leveraging them to underpin neural circuit diagrams, a method which, at last, allows deep learning architectures to be comprehensively and rigorously expressed.

1 Introduction

Category theory enables a rigorous and abstract understanding of composition, offering tools for studying a wide range of systems. One of such tools are diagrams. Cohesive diagrams need to obey strict compositionality rules, making category theory the natural means by which to approach them. In describing systems, diagrams offer a rigorous framework to visualize and analyze them. Furthermore, by diagramming a system, its fundamental abstract structure is revealed.

Categorical diagrams encode the compositional structure of systems, naturally reflecting their algebraic semantics and exchanging the management of symbols for the manipulation of diagrams. The most prominent diagrams in category theory are monoidal string diagrams [17, 14].

These have been applied, or found to underlie, diagrams in various domains. This includes Feynman diagrams in physics, Petri nets for engineered systems, diagrams for quantum computing, and shared theoretical structures in physics, logic, and topology [3, 5, 12]. Monoidal string diagrams offer a common unified approach, letting theorems be applied across domains.

Current diagrammatic approaches, however, are specialized to focus on certain algebraic concepts. They employ topological isotopy rules which ensure that diagrams correspond to topological manipulations which shift around components. This limits them. Monoidal string diagrams, for instance, rely on a privileged product to combine morphisms. This struggles to represent systems where the actions of both linear and non-linear operations are critical to represent, for instance. Furthermore, monoidal string diagrams lack a natural expression for functors and, hence, fail to express many insightful categorical concepts.

Alternative string diagrams [11, 13] address functors and natural transformations but lack the means to readily express products. These diagrams use colors to represent separate categories, and strings are functors showing boundaries between them. Natural transformations are nodes on these boundaries, and may introduce additional functors, acting as intersections between boundaries. These diagrams are focused on advanced categorical concepts, the behaviour of monads and adjunctions, for instance. A variety of existing diagrammatic approaches are showcased in the appendix, Section A.

For applied cases, however, streamlined diagrams are preferred. While morphisms, functors, and natural transformations frequently arise in practice, color coding categories impose distraction but little benefit. Furthermore, expressing both functors and products in either scheme is unclear and obfuscates the underlying expression. Proofs related to areas in which diagrams are not-specialized are difficult to parse. The topological formality of current diagrams, then, can be seen as a hindrance to generalization, cross-domain understanding, and application.

Instead, it will be valuable to distill the construction of diagrams down to clear principles which ensures we can construct a reliable correspondence between categorical expressions and diagrams. Rather than specializing diagrams for a particular case, we will ensure they can robustly represent categorical expressions and will introduce new features as needed. These new features will each provide local graphical intuition and will correspond with algebraic rules, but will not be required to be consistent with an overarching topological regularity. We will prefer to use colors for annotation rather than expression, granting additional utility for explaining and applying diagrams.

This leads to functor string diagrams [2], a principled approach to constructing diagrams for various systems. We will demand that diagrams can be readily corresponded to standard object-morphism-object-etc. expressions. New features will be introduced as needed, with each implementing an algebraic rule in a graphically intuitive manner that motivates its use. Therefore, they can reveal the algebraic structure in a range of cases, from formal proofs to the design of deep learning models, in a clearly understandable manner.

Our Contributions

In this work, the two principles underlying functor string diagrams, vertical section decomposition and equivalent expression for graphical intuition, are introduced. The disciplined diagrams constructed from these principles are shown to provide both theoretical and applied utility. We show this utility by presenting a straightforward proof of the Yoneda lemma, showing existing monoidal string diagrams to be a specific instance of functor string diagrams, and finally show

that they provide mathematical rigor to neural circuit diagrams, addressing the open-problem of comprehensively communicating the details of deep learning models.

2 The Foundations of Functor String Diagrams

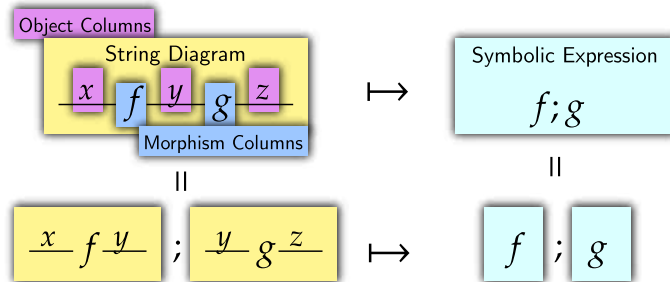
Functor string diagrams fundamentally take the perspective that categorical diagrams should be flexible yet rigorous, adaptable to various use cases while maintaining clear correspondence to underlying symbolic expressions and having features which guide intuition regarding the application of theorems. This is achieved by establishing two principles for the construction of diagrams.

This section will introduce and highlight these principles. If they are followed, we will attain flexible, disciplined diagrams that can represent just about any categorical concept. Compared to monoidal string diagrams which strictly demand isotopy, we have the flexibility to introduce functors and natural transformations in a rigorous, intuitive manner. Furthermore, these diagrams can be adapted to different use cases. As long as we are aware of the equivalent expressions employed, we always have diagrams corresponding to unique up to isomorphism categorical expressions.

2.1 Vertical Section Decomposition

Functor string diagrams can be decomposed into vertical columns. Vertical columns either represent the current object, or a morphism acting on that object. Below in Figure 1, we can see a basic functor string diagram mapping through $f : x \rightarrow y$ and $g : y \rightarrow z$ to construct $f;g : x \rightarrow z$. (*representing forward composition.*) Note how we have vertical columns for the objects x, y, z , and other vertical columns for the morphisms f and g .

Figure 1: Here, we have a morphism with multiple functor wires present. Using the principle of equivalent expression, this can be re-expressed into an easier-to-understand form.



This gives us the first principle for functor string diagrams, vertical section composition, which states that all vertical sections of a diagram correspond to either an object or a morphism. Furthermore, as we express identities by straight lines, the same as objects, the construction of diagrams via vertical section decomposition corresponds with the composition and identity axioms of a category. Therefore, there is a one-to-one correspondence between valid composed functor string diagrams and valid composed symbolic categorical expressions.

2.2 Equivalent Expression

In practice, we will want to introduce additional features to diagrams to represent functors, natural transformations, and other constructs. To maintain cohesion, we will demand that our means of representing these additional features ultimately corresponds to columns of objects and morphisms, abiding by the first principle. This grants the second principle, that all additional features are equivalent expressions of existing features to support graphical intuition.

This guides us towards expressing functors. Functors are morphisms from \mathbf{C} to \mathbf{D} which map objects and morphisms so that composition is preserved. Fundamentally, we require that $f : x \rightarrow y$ maps to $Ff : Fx \rightarrow Fy$ and that $F(f;g) = Ff;Fg$. We express an object Fx by a vertical column with a functor wire F placed above x , and Ff by a vertical column with the functor wire F placed above f , as shown in Figure 2. When constructing $Ff;Fg$, we see that graphical intuition supports its equality to $F(f;g)$.




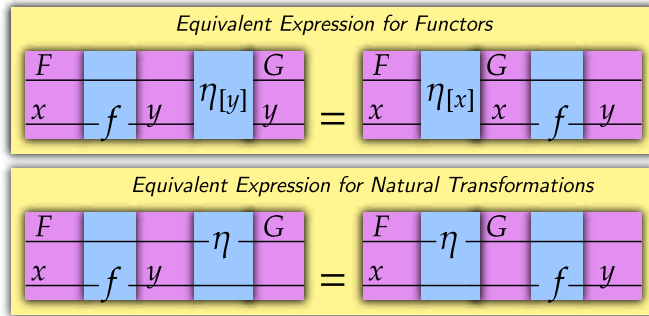
Figure 2: Objects and morphisms in the image of functors can be explicitly expressed. Functors obey the axioms that $f : x \rightarrow y$ maps to $Ff : Fx \rightarrow Fy$ and $F(f;g) = Ff;Fg$. Therefore, we develop an equivalent expression which graphically enforces these axioms, necessarily mapping a morphism along with its objects and necessarily mapping composed morphisms.

Similarly, we find that natural transformations also have an intuitive equivalent expression. We can write the defining axiom of natural transformations at the top, wherein the functors can be equivalently expressed with functor wires. We notice that the natural transformation can move along the functor line, hence, we establish an equivalent expression for natural transformation components by writing them on the functor wire. This is similar to the methods used by Marsden and Nakahira [11, 13] and ensures that the axioms of diagrams reflect the axioms of symbols.

3 The Yoneda Lemma with Functor String Diagrams

To prove the algebraic utility of functor string diagrams, we will be expressing the Yoneda lemma in a natural manner. The Yoneda lemma, a fundamental concept in category theory, has its relevance to various domains potentially obscured when the role of natural transformations is not evident. As traditional monoidal string diagrams can not consider functors or natural transformations, a novel

Figure 3: We develop equivalent expressions to consider natural transformations. We see that natural transformations move along the functor wire and, hence, we show them as morphisms on the functor wire.

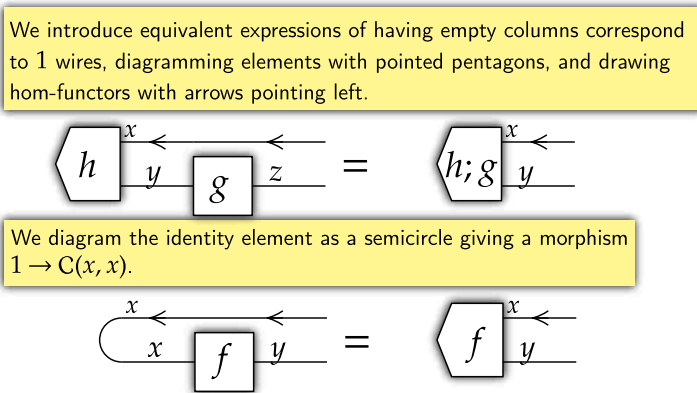


diagrammatic method is needed to understand it. Here, we will provide a proof of the Yoneda lemma, borrowing from Nakahira, but which uses the principles laid out to offer a brief, intuitive proof.

The Yoneda lemma relates to hom-functors and the category **Set**. In **Set**, objects are sets and morphisms are functions between them. In **Set**, the elements of sets correspond to morphisms from the set of one element, 1 , to the object corresponding to the set. In **Set**, morphisms are functions. Functions $f : x \rightarrow y$ are defined by the output element $\varphi; f : 1 \rightarrow y$ produced for each input element $\varphi : 1 \rightarrow x$, and any function pairing each input to any output can be constructed.

For any category **C** where morphisms between objects forms a set (*locally small*), there is a hom-functor for each object x which maps from **C** to **Set**, sending objects y to the *set* of morphisms $\mathbf{C}(x, y)$ and morphisms $g : y \rightarrow z$ to a morphism $\mathbf{C}(x, g) : \mathbf{C}(x, y) \rightarrow \mathbf{C}(x, z)$. The hom-functor morphism corresponds to the following operation in Figure 4.

Figure 4: The hom-functor maps morphisms $g : y \rightarrow z$ to a function between sets $\mathbf{C}(x, g) : \mathbf{C}(x, y) \rightarrow \mathbf{C}(x, z)$. This function implements composition with $h : x \rightarrow y$ when provided with the element of the set $\mathbf{C}(x, y)$ corresponding to h . We can introduce a number of equivalent expressions to more clearly visualize these expressions. This includes not drawing 1 object wires.



Regarding hom-functors, the Yoneda lemma makes the following claim.

Theorem 1 (Yoneda Lemma) *The natural transformations $\eta : \mathbf{C}(x, _) \rightarrow \Phi$, both being functors $\mathbf{C} \rightarrow \mathbf{Set}$, are in one-to-one correspondence with the elements $\varphi \in \Phi_x$, the set the object $x \in \text{Ob}(\mathbf{C})$ is mapped to by Φ :*

$$\text{Nat}(\mathbf{C}(x, _), \Phi) \cong \Phi_x.$$

Proof. We start with an element $\varphi \in \Phi x$. We represent this as a morphism $1 \rightarrow \Phi x$. Diagrammatically, we see that this provides a hook for a morphism $f : x \rightarrow y$ to attach, giving an element $(\Phi f)(\varphi) \in \Phi y$. We will define a function mapping from any morphism $f : x \rightarrow y$ to the outcome of this operation, hence, its a morphism $\varphi_{[y]}^* : \mathbf{C}(x, y) \rightarrow \Phi y$. This morphism $\mathbf{C}(x, y) \rightarrow \Phi y$ exists by virtue of the fact that a function mapping from any set of inputs to outputs is allowed in **Set**.

$$\begin{array}{c} x \\ \leftarrow \\ \text{f} \\ \text{y} \end{array} \begin{array}{c} \Phi \\ \varphi_{[y]}^* \\ \text{y} \end{array} \stackrel{\text{def}}{=} \begin{array}{c} \Phi \\ \varphi \\ \text{x} \end{array} \begin{array}{c} \text{f} \\ \text{y} \end{array}$$

We need to confirm whether this is a natural transformation. We do this by testing the above with a generic function $g : y \rightarrow z$.


$$\begin{array}{c} x \\ \leftarrow \\ \text{f} \\ \text{y} \end{array} \begin{array}{c} \Phi \\ \varphi_{[y]}^* \\ \text{y} \end{array} \begin{array}{c} \text{g} \\ \text{z} \end{array} \stackrel{\Phi}{=} \begin{array}{c} \Phi \\ \varphi \\ \text{x} \end{array} \begin{array}{c} \text{f} \\ \text{y} \end{array} \begin{array}{c} \text{g} \\ \text{z} \end{array} \\ \stackrel{\Phi}{=} \begin{array}{c} \Phi \\ \varphi \\ \text{x} \end{array} \begin{array}{c} \text{f;g} \\ \text{z} \end{array} \\ \stackrel{\Phi}{=} \begin{array}{c} x \\ \leftarrow \\ \text{f;g} \\ \text{z} \end{array} \begin{array}{c} \Phi \\ \varphi_{[z]}^* \\ \text{z} \end{array} \\ \stackrel{\Phi}{=} \begin{array}{c} x \\ \leftarrow \\ \text{f} \\ \text{y} \end{array} \begin{array}{c} \text{g} \\ \text{z} \end{array} \begin{array}{c} \Phi \\ \varphi_{[z]}^* \\ \text{z} \end{array}$$

These functions are equal for all elements in the domain $\mathbf{C}(x, y)$, making them equal.

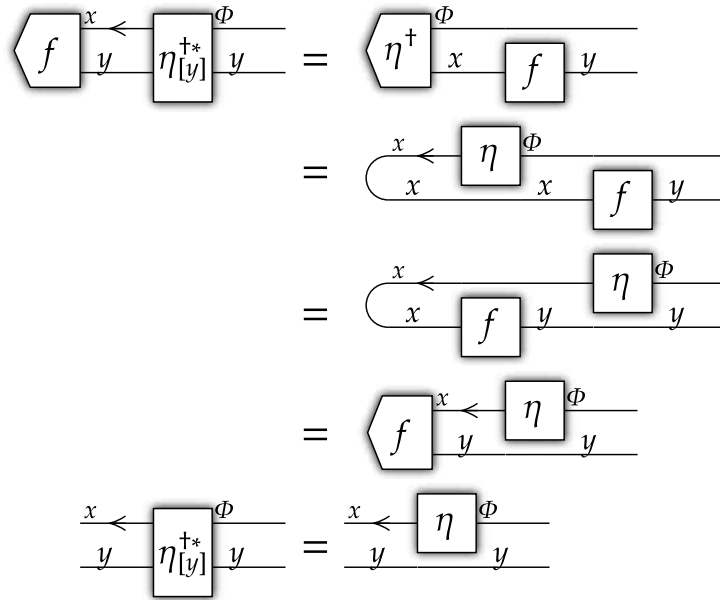
$$\begin{array}{c} x \\ \leftarrow \\ \text{f} \\ \text{y} \end{array} \begin{array}{c} \Phi \\ \varphi_{[y]}^* \\ \text{y} \end{array} \begin{array}{c} \text{g} \\ \text{z} \end{array} \stackrel{\Phi}{=} \begin{array}{c} x \\ \leftarrow \\ \text{y} \end{array} \begin{array}{c} \text{g} \\ \text{z} \end{array} \begin{array}{c} \Phi \\ \varphi_{[z]}^* \\ \text{z} \end{array} \\ \begin{array}{c} x \\ \leftarrow \\ \text{f} \\ \text{y} \end{array} \begin{array}{c} \Phi \\ \varphi^* \\ \text{y} \end{array} \begin{array}{c} \text{g} \\ \text{z} \end{array} \stackrel{\Phi}{=} \begin{array}{c} x \\ \leftarrow \\ \text{y} \end{array} \begin{array}{c} \text{g} \\ \text{z} \end{array} \begin{array}{c} \Phi \\ \varphi^* \\ \text{z} \end{array}$$

Therefore, we have a procedure “*” which maps from elements $\varphi \in \Phi x$ to natural transformations $\varphi^* : \mathbf{C}(x, _) \rightarrow \Phi$. The next phase of the proof is to show that there are elements corresponding to natural transformations, and that taken together, these processes form a bijection.

Starting with a natural transformation $\eta : \mathbf{C}(x, _) \rightarrow \Phi$, we use the identity to obtain an element $1 \rightarrow \Phi x$. This gives a map “†” from natural transforms to elements, $\dagger : \text{Nat}(\mathbf{C}(x, _), \Phi) \rightarrow \Phi x$.



Now, we will find the natural transformation $\eta^{\dagger*}$ corresponding to the element η^\dagger .



This shows that mapping η^\dagger back to a natural transformation recovers it. Therefore, the map from elements to natural transformations is a bijection. This implies there is a one-to-one correspondence between natural transformations and elements, completing the proof. ■

The Yoneda lemma, fundamentally, says nothing more or less than that the following graphical move shown in Figure 5 is allowed. Elements can become natural transformation components, and vice-versa.




Figure 5: The Yoneda lemma stated with our diagrammatic approach. It shows that elements $\varphi_\eta \in \Phi x$ correspond to natural transformations $\eta_\varphi : \mathbf{C}(x, _) \rightarrow \Phi$. This is equivalent to the graphically intuitive rule where elements can “jump up” and become natural transformations, and vice-versa.


4 Functor String Diagrams, Products, and Monoidal String Diagrams

In addition to having utility in streamlining algebraic proofs and understanding, functor string diagrams encompass monoidal string diagrams as a specific case. Hence, they are shown to capture existing diagrams and provide a strict improvement over current methods which have found success representing Feynman diagrams, Petri nets, and other structures mentioned in the introduction.

Monoidal string diagrams rely on a category \mathbf{C} having a functor \otimes from $\mathbf{C} \times \mathbf{C} \rightarrow \mathbf{C}$. $\mathbf{C} \times \mathbf{D}$ is the product category wherein objects are tuples and morphisms are tuples from \mathbf{C} and \mathbf{D} . We equivalently express objects and morphisms in a product category by a double-dashed line separating the constituent objects or morphisms as in Figure 6.

Figure 6: Product categories such as $\mathbf{C} \times \mathbf{D}$ are represented by a double-dashed line. This can represent the combination of expressions from any two categories.

We develop double-dashed lines as an equivalent expression to represent product categories $\mathbf{C} \times \mathbf{D}$, where objects and morphisms are tuples.



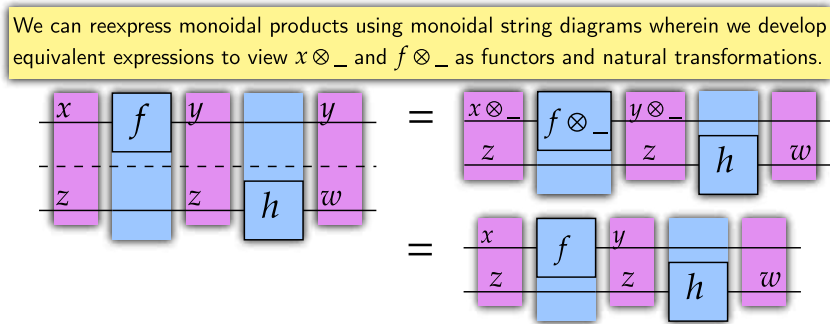
Therefore, $\mathbf{C} \times \mathbf{C}$ is represented as two expressions from the same category separated by a double dashed line. A functor applied onto a product expression needs to be indicated as distinct from a functor on the upmost expression, and this is achieved with a bold line. Thereby, $\otimes : \mathbf{C} \times \mathbf{C} \rightarrow \mathbf{C}$ can be represented with a bold wire labeled “ \otimes ” (see appendix Figure 19). Furthermore, monoidal categories typically have a privileged product. In this case, it is represented with a single dashed line separating expressions.

Second, we can represent monoidal categories without a dashed line. This is done by realizing that a monoidal product $\otimes : \mathbf{C} \times \mathbf{C} \rightarrow \mathbf{C}$ grants every object in $x \in \text{Ob}(\mathbf{C})$ a functor $x \otimes _ : \mathbf{C} \rightarrow \mathbf{C}$ and every morphism $f : x \rightarrow y$ a natural transformation $f \otimes _ : x \otimes _ \rightarrow y \otimes _$. This fact can be observed from the diagram where f can naturally move over h , with components $f \otimes \text{Id}_z^z$ and $f \otimes \text{Id}_w^w$. This is shown in Figure 7. Monoidal string diagrams, therefore, are functor string diagrams where $x \otimes _$ is equivalently expressed by a wire labeled x , and natural transformations $f \otimes _$ are expressed with boxes labeled f on functor wires $x \otimes _$ to $y \otimes _$.

If we look at how monoidal string diagrams are defined, we see that they obey vertical section decomposition and equivalent expression. A survey by Selinger [17] represents a standard introduction. It introduces the features of monoidal string diagrams as re-expressions of symbolic categories, as seen in appendix Figure 19.

Monoidal string diagrams focus on powerful algebraic correspondence between diagrammatic and categorical expressions. This *coherence* property states that well-formed diagrams exactly correspond to well-formed expressions. This property is so powerful that it allows all the details of monoidal categories, isomorphisms, associators, and so forth, to follow from the axioms of the graphical language and planar isotopy. Monoidal string diagrams, then, can be used to define

Figure 7: By a certain choice of equivalent expression, wherein we choose to view monoidal products as functors and natural transformations, one can see that monoidal string diagrams are a certain specification of functor string diagrams.



what monoidal categories are and hence offer an extremely powerful perspective from which to understand them.

This level of correspondence is not the focus of functor string diagrams. Demanding this isotopy raises a barrier to diagramming and understanding additional structures. To the degree that graphical intuition is useful, it is far simpler to define properties by their local interactions. For instance, natural transformations acting on functor wires follows from a simple local definition. This approach, which does not demand isotopy, allows more features to be introduced with each having some appropriate form of graphical intuition. The graphical intuition embedded into equivalent expressions can coincidentally serve as axioms for the feature, however, this is not the focus of functor string diagrams.

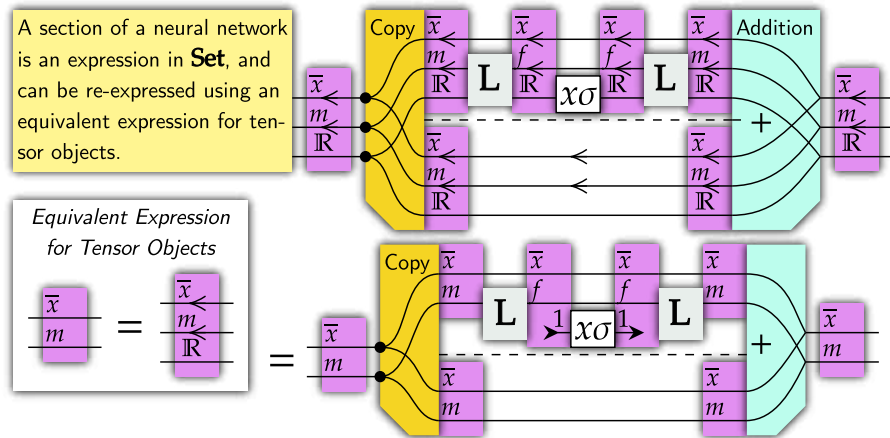
Furthermore, functor string diagrams are sufficiently coherent. All categorical expressions are, ultimately, a chain of morphisms acting on objects, which vertical section decomposition ensures we can perfectly represent. When certain expressions are isomorphic, it is often straightforward to encode that isomorphism as an equivalent expression as we ultimately demand that diagrams correspond to a unique categorical expression up to isomorphism.

5 Neural Circuit Diagrams

Deep learning systems are on the forefront of human technology, with applications in a variety of fields ranging from the scientific investigation of protein folding [18, 4] to immense commercial developments. Currently, ad-hoc methods are used to explain these systems, limiting our understanding and obfuscating implementation. Existing categorical machine learning, and theoretical machine learning more generally, often focus on toy models which do not reflect those used in practice given the limit tools to express the details of models [16, 7, 21, 19] (see also appendix Figure 18). This stems from a fundamental limitation of monoidal string diagrams, as they are unable to express both the details of separate data and the interaction of axes [6]. Functor string diagrams, then, have an exciting application regarding formalizing neural circuit diagrams [1, 2], a categorical method for comprehensively communicating deep learning models.

The Form of Deep Learning Models. Deep learning models are functions between data types, meaning they are expressions in **Set**. Typically, data types are tensors of an underlying data type \mathbb{R} . Therefore, objects are almost always of hom-functors applied to \mathbb{R} . We develop an equivalent expression, whereby wires without arrows represent those objects applied as hom-functors to \mathbb{R} . This allows broadcasted functions between tensors to be represented as in Figure 8. Furthermore, to represent functions $\mathbb{R} \rightarrow \mathbb{R}$, we exploit that the 1-hom-functor is an isomorphism, letting it be freely introduced and removed (unsqueezing and squeezing), which we introduce and remove with a wire with an arrowhead.

Figure 8: For deep learning models, objects typically consist of hom-functors to \mathbb{R} . We develop an equivalent expression which lets this be readily expressed.

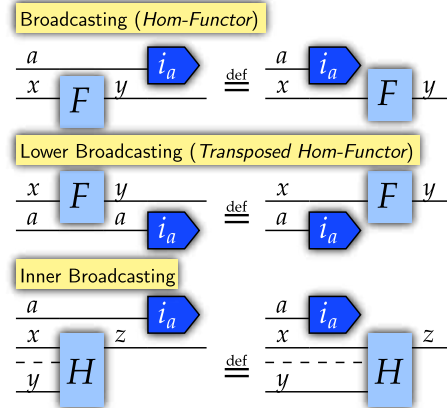


Understanding Expressions with Indexes. As we are operating in **Set**, every object a has a set of coprojections given by its elements $\langle i_a | : 1 \rightarrow a$. Morphisms $a \rightarrow b$ can be uniquely identified and freely constructed from pre-compositions with elements. By the Yoneda lemma, natural transformations $\mathbf{Set}(a, _) \rightarrow \mathbf{Set}(1, _)$ correspond to elements of $\mathbf{Set}(1, a)$. Therefore, elements $\langle i_a | : 1 \rightarrow a$ become natural transformations $|i_a^\leftarrow\rangle_x^a : \mathbf{Set}(a, x) \rightarrow \mathbf{Set}(1, x)$ which we call *indexes*. They form a set of projections (see appendix Figure 22 and 23), letting morphisms to $\mathbf{Set}(a, x)$ be uniquely identified and freely constructed from post-composition with indexes.

Broadcasting. Using indexes, we can define and construct various useful forms of functions. This includes broadcasting, corresponding to hom-functors, and inner broadcasting, corresponding to broadcasting within tuple segments. These expressions use continuous wires to indicate the naturality of indexes, as in Figure 9.

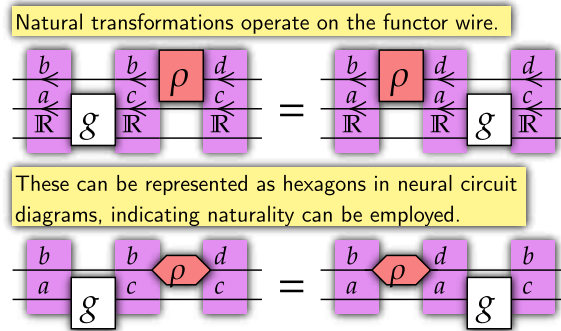
Natural Transformations. By the Yoneda lemma, functions on sets $x \rightarrow y$ correspond to functions manipulating the indexes of tensors $\mathbb{R}^y \rightarrow \mathbb{R}^x$. We represent natural transformations in neural circuit diagrams using hexagons (see Figure 10), and note that they travel along continuous wires

Figure 9: Continuous wires implement naturality, letting natural transformations, including indexes, be transmitted, defining these expressions.



the same way as indexes (see Figure 11). Continuous wires defined by naturality with indexes, then, can be seen to grant naturality to natural transformations more generally.

Figure 10: Natural transformations can be shown with hexagons, and they have naturality regarding hom-functor broadcasting.




Einops. Einstein operations refers to a common class of rearranging and contracting (summation) over axes. This covers a common class of operations, the details of which are critical yet challenging to express with traditional methods. Einstein operations can be expressed by associating desired input with output indexes symbolically, such as with Einstein tensor notation or the Einops package [15]. With neural circuit diagrams, they correspond to continuous wiring associating input and output axes. These operations are linear and often require an outer product to first be taken (element-wise multiplication), which we diagram by terminating a single dotted line.

Scaled dot-product attention. Using these tools, we can express a significant yet complicated deep learning algorithm, multi-head attention, clearly. This algorithm requires interaction between specific axes and independent data, which is hard to express with traditional methods (see appendix Figure 24 for a typical presentation). With neural circuit diagrams, the specific interactions for scaled dot-product attention (SDPA) can be clearly seen as in Figure 13. Furthermore, multi-head

Figure 11: By testing this inner broadcasting expression with indexes k_b , we see that the naturality of indexes as indicated in Figure 9 implies a naturality of natural transformations along continuous wires.

As $\langle k_b |; \rho$ is an index, it is natural along continuous wires. This lets us show that natural transformations are also natural along continuous wires.



The termination of a dashed line indicates an outer product. Over \mathbb{R}^1 , this is just multiplication.




Figure 12: Natural operations, including indexes, travel along continuous wires. This means the above expression, for some choice of indexes $i_{\bar{x}}$, $j_{\bar{x}}$, k_h placed after the Einstein operation, can be seen to be a contraction over the indexes placed before the operation.

attention which performs SDPA h -times and aggregates the result, can be shown be shown with neural circuit diagrams in Figure 14.

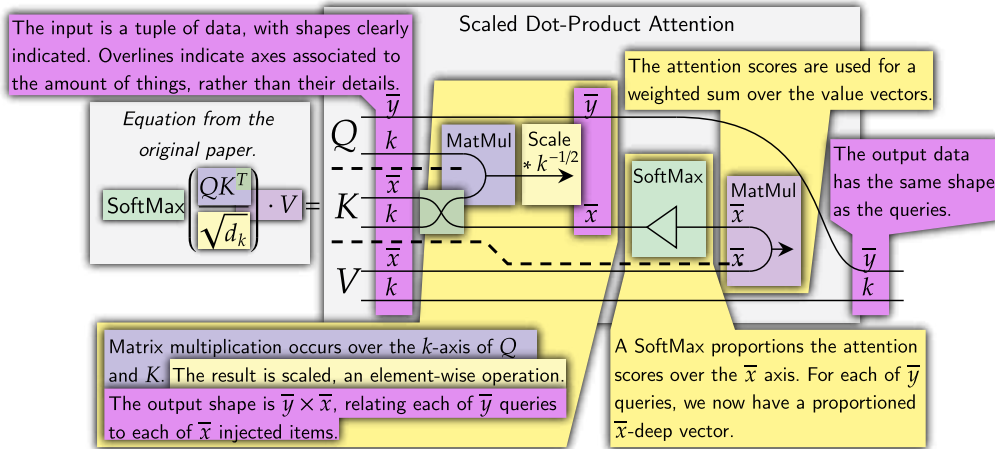


Figure 13: Scaled dot-product attention can be far more clearly expressed with neural circuit diagrams. The interaction between axes and the changing shape of data is naturally shown by diagrams.




Figure 14: Conceptually, multi-head attention involved broadcasting SDPA over h parallel operations and aggregating the result. The queries, keys, and values are produced and recombined by adding an additional axis to the learned linear layers.

6 Conclusion

In this work, we stated the need for a general means of diagrammatically expressing categorical ideas. Functor string diagrams are a general and principled approach to developing categorical diagrams that have robust mathematics for theoretical investigations, encompass existing monoidal string diagram, and have clear utility for application. Compared to the existing paradigm of monoidal string diagrams and string diagrams, functor string diagrams offer a flexible yet principled approach that ensures complex ideas can be clearly expressed. Their ability to express both functors and products allows for theoretical ideas to be clearly understood, and allows for applied systems to be straightforwardly modeled.

References

- [1] Vincent Abbott (2023): *Neural Circuit Diagrams: Robust Diagrams for the Communication, Implementation, and Analysis of Deep Learning Architectures*. *Transactions on Machine Learning Research*. Available at <https://openreview.net/forum?id=RyZB4qXEgt>.
- [2] Vincent Abbott (2023): *Robust Diagrams for Deep Learning Architectures: Applications and Theory*. Honours Thesis, The Australian National University, Canberra.
- [3] Samson Abramsky & Bob Coecke (2008): *Categorical quantum mechanics*, doi:10.48550/arXiv.0808.1023. Available at <http://arxiv.org/abs/0808.1023>. ArXiv:0808.1023 [quant-ph].
- [4] Jumper et al. (2021): *Highly accurate protein structure prediction with AlphaFold*. *Nature* 596(7873), pp. 583–589, doi:10.1038/s41586-021-03819-2. Available at <https://www.nature.com/articles/s41586-021-03819-2>. Number: 7873 Publisher: Nature Publishing Group.
- [5] John C. Baez & Mike Stay (2010): *Physics, Topology, Logic and Computation: A Rosetta Stone*. 813, pp. 95–172, doi:10.1007/978-3-642-12821-9-2. Available at <http://arxiv.org/abs/0903.0340>. ArXiv:0903.0340 [quant-ph].
- [6] David Chiang, Alexander M. Rush & Boaz Barak (2022): *Named Tensor Notation*. *Transactions on Machine Learning Research*. Available at <https://openreview.net/forum?id=hVT7SHlilx>.
- [7] Robin Cockett, Geoffrey Cruttwell, Jonathan Gallagher, Jean-Simon Pacaud Lemay, Benjamin MacAdam, Gordon Plotkin & Dorette Pronk (2019): *Reverse derivative categories*, doi:10.48550/arXiv.1910.07065. Available at <http://arxiv.org/abs/1910.07065>. ArXiv:1910.07065 [cs, math].
- [8] G. S. H. Cruttwell, Bruno Gavranović, Neil Ghani, Paul Wilson & Fabio Zanasi (2021): *Categorical Foundations of Gradient-Based Learning*, doi:10.48550/arXiv.2103.01931. Available at <http://arxiv.org/abs/2103.01931>. ArXiv:2103.01931 [cs, math].
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren & Jian Sun (2015): *Deep Residual Learning for Image Recognition*. CoRR abs/1512.03385. Available at <http://arxiv.org/abs/1512.03385>. ArXiv:1512.03385.
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren & Jian Sun (2016): *Identity Mappings in Deep Residual Networks*. In Bastian Leibe, Jiri Matas, Nicu Sebe & Max Welling, editors: *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part IV, Lecture Notes in Computer Science* 9908, Springer, pp. 630–645, doi:10.1007/978-3-319-46493-0_38. Available at https://doi.org/10.1007/978-3-319-46493-0_38.
- [11] Dan Marsden (2014): *Category Theory Using String Diagrams*. CoRR abs/1401.7220. Available at <http://arxiv.org/abs/1401.7220>. ArXiv:1401.7220.
- [12] Jose Meseguer & Ugo Montanari (1990): *Petri nets are monoids*. *Information and Computation* 88(2), pp. 105–155, doi:10.1016/0890-5401(90)90013-8. Available at <https://www.sciencedirect.com/science/article/pii/0890540190900138>.
- [13] Kenji Nakahira (2023): *Diagrammatic category theory*, doi:10.48550/arXiv.2307.08891. Available at <http://arxiv.org/abs/2307.08891>. ArXiv:2307.08891 [math].
- [14] Robin Piedeleu & Fabio Zanasi (2023): *An Introduction to String Diagrams for Computer Scientists*, doi:10.48550/arXiv.2305.08768. Available at <http://arxiv.org/abs/2305.08768>. ArXiv:2305.08768 [cs].

- [15] Alex Rogozhnikov (2021): *Einops: Clear and Reliable Tensor Manipulations with Einstein-like Notation*. Available at <https://openreview.net/forum?id=oapKSVM2bcj>.
- [16] Andrew M. Saxe, James L. McClelland & Surya Ganguli (2014): *Exact solutions to the nonlinear dynamics of learning in deep linear neural networks*. In Yoshua Bengio & Yann LeCun, editors: *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*. Available at <http://arxiv.org/abs/1312.6120>.
- [17] Peter Selinger (2009): *A survey of graphical languages for monoidal categories*, doi:10.1007/978-3-642-12821-9-4. Available at <https://arxiv.org/abs/0908.3347v1>.
- [18] Andrew W. et al. Senior (2020): *Improved protein structure prediction using potentials from deep learning*. *Nature* 577(7792), pp. 706–710, doi:10.1038/s41586-019-1923-7. Available at <https://www.nature.com/articles/s41586-019-1923-7>. Number: 7792 Publisher: Nature Publishing Group.
- [19] Dan Shiebler, Bruno Gavranovic & Paul W. Wilson (2021): *Category Theory in Machine Learning*. CoRR abs/2106.07032. Available at <https://arxiv.org/abs/2106.07032>. ArXiv: 2106.07032.
- [20] Ashish et al. Vaswani (2017): *Attention is all you need*. In: *Advances in neural information processing systems*, 30, Curran Associates, Inc. Available at <https://proceedings.neurips.cc/paper-files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf>.
- [21] Paul Wilson & Fabio Zanasi (2022): *Categories of Differentiable Polynomial Circuits for Machine Learning*, doi:10.48550/arXiv.2203.06430. Available at <http://arxiv.org/abs/2203.06430>. ArXiv:2203.06430 [cs, math].

Appendix

A Current Categorical Diagrams

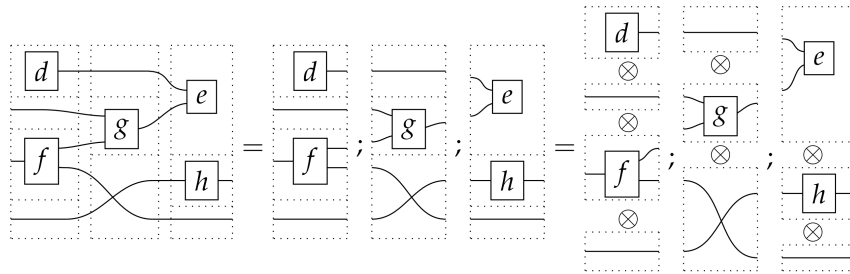


Figure 15: A traditional monoidal string diagram from page 18 of [14]. This expression clearly shows the objects and morphisms at each stage, and has graphical isotopy whereby topological transforms of diagrams lead to valid, isomorphic expressions. However, it does not show the details of axes should the objects be Cartesian tuples. Furthermore, it cannot show functors which are not monoidal products of objects nor natural transformations which are not morphisms between the monoidal product of objects. The useful properties of monoidal string diagrams are encompassed by functor string diagrams, as shown in Section 4.




Figure 16: A string diagram from Marsden [11]. These expressions use coloured regions to show the current category, and display functors as left-to-right boundaries between coloured regions. Natural transformations are dots on functors or their intersection. These can clearly display the action of functors, natural transformations, and more complicated structures such as adjunctions, monads, and (co)limits, and display graphical isotopy, implementing the axioms of these structures graphically. However, they fail to be clearly decomposable for applied cases, and have a cumbersome representation of products. The behaviour of functors and natural transformations using functor string diagrams are shown in Section 2.




Figure 17: Nakahira's diagrams [13] extends Marsden's diagrams to consider substitutions and various algebraic proofs, including the Yoneda lemma shown above. These diagrams have powerful tools for the consideration of products, bifunctors, and other constructs. However, streamlined simplicity is preferred for practical application and intuitive understanding, which is shown in the functor string diagram proof for the Yoneda lemma in Section 3.




Figure 18: Categorical machine learning such as from [8] uses monoidal string diagrams with Cartesian products, along with the **Para** construct which allows parameters to be introduced from above. These do not express the details of axes, and hence focus on toy models rather than those used in practice. Since 2015 [9, 10], for example, deep learning has employed residual connections which are significantly different from linear-bias-activation toy models. Functor string diagrams underly neural circuit diagrams (see Section 5), which can rigorously and comprehensively express deep learning models.

B Equivalent Expression of Products

With functor string diagrams, we can express bifunctors $\mathbf{C} \times \mathbf{C} \rightarrow \mathbf{C}$ as bolded functor wires over an expression in $\mathbf{C} \times \mathbf{C}$, as in Figure 19. As we can see, functor string diagrams provide a flexible framework to understand the connections between various means of viewing expressions.

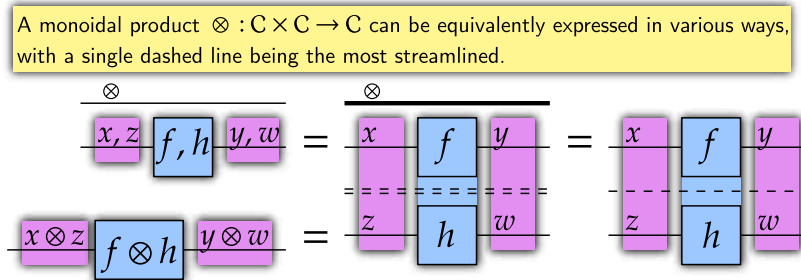


Figure 19: These equivalent expressions for the diagramming of product categories and monoidal products offer immense utility. We see how different expressions are related and can freely adopt whichever approach is the most useful.


C Typical Introduction to Monoidal String Diagrams

Typical introductions to monoidal string diagrams introduce the fundamental structures by elaborating how objects and morphisms are represented. Below, tables 1 and 2 from [17] are shown in Figure 20 and 21. This lets us see that monoidal string diagrams are, fundamentally, diagrams with vertical sections corresponding to objects and morphisms equipped with equivalent expressions.

Figure 20: Table 1 from [17], “The graphical language of categories”.

Object	A	\xrightarrow{A}
Morphism	$f : A \rightarrow B$	$\xrightarrow{A} \boxed{f} \xrightarrow{B}$
Identity	$\text{id}_A : A \rightarrow A$	\xrightarrow{A}
Composition	$t \circ s$	$\xrightarrow{A} \boxed{s} \xrightarrow{B} \boxed{t} \xrightarrow{C}$

Figure 21: Table 2 from [17], “The graphical language of monoidal categories”.



D Elements, Indexes, and Projections

Figure 22: Elements of a set $1 \rightarrow a$ form a set of coprojections which can uniquely identify and construct functions $a \rightarrow _$.

By the Yoneda lemma, we have a correspondence between elements $i_a : 1 \rightarrow a$ and natural transforms $C(a, _) \rightarrow C(1, _)$.



In **Set**, we have a natural isomorphism $C(1, _) \cong _$ letting us equivalently express $C(1, _)$ without drawing the functor wire.




Figure 23: The associated natural transformations $\mathbf{Set}(a, _) \rightarrow _$, for a set of projections, letting functions $_ \rightarrow \mathbf{Set}(a, _)$ be uniquely identified.

Natural transformations associated to elements, called indexes, uniquely identify morphisms. If two morphisms $1 \rightarrow \mathbf{Set}(a, x)$ match for all indexes, their underlying morphisms match for all elements. Hence, they are equal, and we observe that indexes for projections.

$$\begin{aligned}
 \forall i_a : 1 \rightarrow a. \quad & \langle f \mid_x \rangle \xrightarrow{a} \langle i_a \mid_x \rangle = \langle g \mid_x \rangle \xrightarrow{a} \langle i_a \mid_x \rangle \\
 \Leftrightarrow \forall i_a : 1 \rightarrow a. \quad & \langle i_a \mid_a \rangle \langle f \mid_x \rangle = \langle i_a \mid_a \rangle \langle g \mid_x \rangle \\
 \Leftrightarrow & \langle f \mid_x \rangle = \langle g \mid_x \rangle
 \end{aligned}$$

E Typical Presentation of Attention

A typical presentation of scaled dot-product attention can be drawn from [20] are shown in Figure 24. This work introduced transformer models. However, its diagrams are difficult to understand, especially with regards to the interaction of axes and the shape of data throughout the model. As can be seen, Figure 13 and 14 are a significant improvement.




Figure 24: The prevailing ad-hoc diagrams used to represent scaled dot-product attention and multi-head attention obfuscate the interaction of axes and how the shape of data changes across operations, making understanding and implementation difficult.