Sheaves on Quivers via a Grothendieck Topology on the Path Category

Eric M. Schmid Jr., Fernando Tohmé, William Chin

Abstract

We construct Grothendieck topologies on the path category of a finite graph, examining both coarse and discrete cases that offer different perspectives on quiver representations. The coarse topology declares each vertex covered by all incoming morphisms, giving the minimal non-trivial Grothendieck topology where sheaves correspond to dual representations via dualization. The discrete topology is the finest possible, forcing sheaves to be locally constant with isomorphic restriction maps. We verify these satisfy Grothendieck's axioms, characterize their sheaf categories, and establish functorial relationships between them. Sheaves on the coarse site arise naturally from quiver representations through dualization, while discrete sheaves correspond to representations of the groupoid completion. This work suggests intermediate topologies could capture subtler representation-theoretic phenomena.

Contents

Introduction

Grothendieck's theory of sites extends topological covering to abstract categories, specifying which families of morphisms $\{U_i \to U\}$ cover an object U according to axioms mimicking topological covers. This framework has proved essential across mathematics, from étale topology in algebraic geometry to canonical topologies in topos theory.

We apply these ideas to path categories of finite directed graphs. Quivers play a central role in representation theory as organizing structures for algebras and modules, making their categorical analysis particularly valuable. The path category $\mathcal{C}_Q$ of a quiver Q captures the graph's combinatorial structure while providing a natural setting for representations.

Given a finite graph Q with vertices Q_0 and directed edges Q_1, where each edge $e \in Q_1$ has source $s(e) \in Q_0$ and target $t(e) \in Q_0$, we assume no loops $e : v \to v$ to avoid pathologies. The path category $\mathcal{C}_Q$ has vertices as objects and directed paths as morphisms, with composition by path concatenation.

We study two extremal Grothendieck topologies on $\mathcal{C}_Q$:

The Coarse Topology: For each vertex v, we declare the maximal sieve $R_v = \mathrm{Hom}_{\mathcal{C}_Q}(-, v)$ covering. This gives the coarsest non-trivial topology where "every way of reaching a vertex covers it." Despite seeming permissive, this yields tractable sheaf conditions and connects directly to quiver representation theory.

The Discrete Topology: Every sieve is covering, giving the finest possible topology. This imposes maximal constraints: sheaves must have isomorphic restriction maps, becoming locally constant.

The contrast between these topologies reveals different aspects of representation theory. Under the coarse topology, quiver representations $V : \mathcal{C}_Q \to \mathbf{Vect}_k$ dualize to sheaves V^* automatically.

The discrete topology forces sheaves to correspond to representations of the groupoid completion, where all morphisms are invertible.

We establish functorial relationships between these sheaf categories, including a natural inclusion from discrete to coarse sheaves with left adjoint. This suggests intermediate topologies could capture more nuanced representation-theoretic phenomena - edge-generated topologies distinguishing actual graph edges, or graded topologies respecting path length.

Our framework places quiver representations within sheaf theory, accessing tools from topos theory while potentially revealing new structural insights. This aligns with contemporary categorical approaches that unify diverse mathematical phenomena.

The paper proceeds by establishing both topologies and verifying Grothendieck's axioms, characterizing their sheaf categories, exploring functorial relationships, and outlining future directions including intermediate topologies and applications to representation theory.

Notation. We fix a base field k. All vector spaces are over k, and $\mathbf{Vect}_k$ denotes the category of finite-dimensional k-vector spaces. For a set S, we write $|S|$ for its cardinality. Composition of morphisms $f : X \to Y$ and $g : Y \to Z$ in a category is denoted $g \circ f$. Identities are denoted id_X.

1 A Grothendieck Topology on the Path Category of a Graph

1.1 The path category of a finite graph

Let $Q = (Q_0, Q_1)$ be a finite directed graph (quiver) with vertex set Q_0 and directed edge set Q_1. For each edge $e \in Q_1$, denote its source by $s(e) \in Q_0$ and its target by $t(e) \in Q_0$. We assume $s(e) \neq t(e)$ for each edge (no loops) to avoid pathologies in the topology definition. We do allow multiple distinct edges between the same pair of vertices.

From Q we form a small category $\mathcal{C}_Q$, called the *path category* or free category on the quiver Q. The objects of $\mathcal{C}_Q$ are the vertices Q_0. For any two vertices $x, y \in Q_0$, a morphism $\alpha : x \to y$ in $\mathcal{C}_Q$ is defined to be a directed path in the quiver from x to y. Concretely, a path α of length $n \geq 1$ is a finite composable sequence of n edges,

$$\alpha = (x = v_0 \xrightarrow{e_1} v_1 \xrightarrow{e_2} v_2 \xrightarrow{e_3} \cdots \xrightarrow{e_n} v_n = y),$$

where each $e_i \in Q_1$ with $s(e_1) = v_0$, $t(e_i) = s(e_{i+1})$ for $1 \leq i < n$, and $t(e_n) = v_n$. We also include a path of length 0 at each vertex v, which serves as the identity morphism $\mathrm{id}_v : v \to v$. Composition in $\mathcal{C}_Q$ is by concatenation of paths: if $\alpha : x \to y$ and $\beta : y \to z$ are paths, then $\beta \circ \alpha : x \to z$ is the path obtained by first traversing α and then β in sequence.

Example 1. *Consider a simple graph with three vertices a, b, c and two edges $e_1 : a \to b$ and $e_2 : b \to c$. The path category $\mathcal{C}_Q$ has objects $\{a, b, c\}$ and morphisms include:*

- *Identity morphisms:* $\mathrm{id}_a, \mathrm{id}_b, \mathrm{id}_c$
- *Edge morphisms:* $e_1 : a \to b$, $e_2 : b \to c$

- *Composed path:* $e_2 \circ e_1 : a \to c$

1.2 Definition of the topology via covering sieves

We now define a Grothendieck topology on $\mathcal{C}_Q$. Recall that a Grothendieck topology on a category $\mathcal{C}$ assigns to each object X a collection $J(X)$ of covering sieves on X satisfying three axioms:

(GT1) The maximal sieve $\mathrm{Hom}(-, X)$ is in $J(X)$

(GT2) (Stability) If $S \in J(X)$ and $f : Y \to X$ is any morphism, then the pullback sieve $f^*(S) \in J(Y)$

(GT3) (Transitivity) If $S \in J(X)$ and R is any sieve on X such that for all $f : Y \to X$ in S we have $f^*(R) \in J(Y)$, then $R \in J(X)$

Definition 1 (Covering sieves)**.** *For each object (vertex) $v \in Q_0$, we declare the maximal sieve*

$$R_v := \mathrm{Hom}_{\mathcal{C}_Q}(-, v) = \{f : u \to v \mid u \in Q_0\}$$

to be a covering sieve of v. Note that R_v consists of all *morphisms with codomain v, not just those factoring through edges. No other sieves are declared as covering except those containing a covering sieve. We denote by J the collection of all covering sieves generated by this specification.*

Remark 1. *This is the coarsest non-trivial topology on $\mathcal{C}_Q$. It differs from the canonical topology, where coverings would consist of jointly epimorphic families. Here, we take the entire hom-set as the cover, which simplifies the sheaf condition considerably.*

1.3 Verification of the Grothendieck topology axioms

Theorem 1. *The assignment $v \mapsto R_v = \mathrm{Hom}_{\mathcal{C}_Q}(-, v)$ defines a Grothendieck topology J on $\mathcal{C}_Q$.*

Proof. We verify the axioms (GT1)–(GT3):
(GT1) Existence of a maximal cover: By definition, $R_v = \mathrm{Hom}(-, v)$ is the maximal sieve on v, and we declared it to be a covering sieve. Hence (GT1) is satisfied.
(GT2) Stability under pullback: Let $S = R_X = \mathrm{Hom}(-, X)$ be a covering sieve on X and $f : Y \to X$ any morphism. The pullback sieve is:

$$f^*(R_X) = \{g : Z \to Y \mid f \circ g \in R_X\} = \{g : Z \to Y \mid f \circ g : Z \to X\} = \mathrm{Hom}(-, Y) = R_Y.$$

Since R_Y is declared to be a covering sieve on Y, we have $f^*(R_X) \in J(Y)$.
(GT3) Transitivity: Let $S = R_X$ be a covering sieve on X, and R any sieve on X such that for all $h : Y \to X$ in S, we have $h^*(R) \in J(Y)$. We must show $R \in J(X)$. Since $R_X = \mathrm{Hom}(-, X)$, we have $\mathrm{id}_X \in R_X$. By hypothesis, $\mathrm{id}_X^*(R) \in J(X)$. But

$$\mathrm{id}_X^*(R) = \{g : Z \to X \mid \mathrm{id}_X \circ g \in R\} = \{g : Z \to X \mid g \in R\} = R.$$

Therefore $R \in J(X)$. $\square$

2 Sheaves of Vector Spaces on the Graph Site

2.1 Presheaves and sheaves on $(\mathcal{C}_Q, J)$

Having established the site, we turn to defining sheaves valued in $\mathbf{Vect}_k$. We begin by recalling the basic definitions and then examine how our particular topology simplifies the sheaf condition.

Definition 2 (Presheaf of vector spaces)**.** *A* presheaf of k-vector spaces *on $\mathcal{C}_Q$ is a contravariant functor $F : \mathcal{C}_Q^{op} \to \mathbf{Vect}_k$. Concretely, this consists of:*

- *For each vertex $v \in Q_0$, a k-vector space $F(v)$*
- *For each morphism $\alpha : u \to v$ in $\mathcal{C}_Q$, a k-linear map $F(\alpha) : F(v) \to F(u)$ (called a restriction map)*

subject to the conditions:

- $F(\mathrm{id}_v) = \mathrm{id}_{F(v)}$ *for all $v \in Q_0$*
- $F(\beta \circ \alpha) = F(\alpha) \circ F(\beta)$ *for composable morphisms $\alpha : u \to v$ and $\beta : v \to w$*

Example 2. *Let Q be the graph with vertices $\{a, b, c\}$ and edges $e_1 : a \to b$, $e_2 : b \to c$. A presheaf F on $\mathcal{C}_Q$ assigns:*

- *Vector spaces: $F(a)$, $F(b)$, $F(c)$*
- *Linear maps: $F(e_1) : F(b) \to F(a)$, $F(e_2) : F(c) \to F(b)$, $F(e_2 \circ e_1) : F(c) \to F(a)$*
- *Identity maps: $F(\mathrm{id}_a)$, $F(\mathrm{id}_b)$, $F(\mathrm{id}_c)$*

with the compatibility condition $F(e_2 \circ e_1) = F(e_1) \circ F(e_2)$.

Definition 3 (Sheaf condition)**.** *A presheaf F on the site $(\mathcal{C}_Q, J)$ is a* sheaf *if for every vertex v and every covering sieve $R_v \in J(v)$, the following diagram is an equalizer:*

$$F(v) \xrightarrow{\epsilon} \prod_{(f:u\to v)\in R_v} F(u) \rightrightarrows \prod_{\substack{(f:u\to v)\in R_v \\ (g:w\to u)}} F(w)$$

where:

- *The map ϵ is given by $\epsilon(s) = (F(f)(s))_{f\in R_v}$*
- *The two maps to the double product are given by:*

$$p_1((s_f)_{f\in R_v}) = (s_{f\circ g})_{f,g} \tag{1}$$
$$p_2((s_f)_{f\in R_v}) = (F(g)(s_f))_{f,g} \tag{2}$$

Since our covering sieve $R_v = \mathrm{Hom}_{\mathcal{C}_Q}(-, v)$ consists of all morphisms into v, we can make this condition more explicit:

Theorem 2 (Simplified sheaf condition)**.** *A presheaf F on $(\mathcal{C}_Q, J)$ is a sheaf if and only if for each vertex v, given any family of vectors $\{s_f \in F(u)\}_{f:u\to v}$ such that for all composable pairs $f : u \to v$ and $g : w \to u$,*

$$F(g)(s_f) = s_{f\circ g}$$

there exists a unique vector $s \in F(v)$ such that $F(f)(s) = s_f$ for all $f : u \to v$.

Proof. Recall that a presheaf F on $(\mathcal{C}_Q, J)$ is a sheaf if and only if for every object v and every covering sieve $R_v \in J(v)$, the following diagram is an equalizer:

$$F(v) \xrightarrow{\epsilon} \prod_{(f:u\to v)\in R_v} F(u) \rightrightarrows \prod_{\substack{(f:u\to v)\in R_v \\ (g:w\to u)}} F(w)$$

where the map ϵ is given by $\epsilon(s) = (F(f)(s))_{f\in R_v}$ and the two parallel morphisms p_1, p_2 are given by:

$$p_1((s_f)_{f\in R_v}) = (s_{f\circ g})_{f,g} \tag{3}$$

$$p_2((s_f)_{f\in R_v}) = (F(g)(s_f))_{f,g} \tag{4}$$

In our specific case, the covering sieve $R_v = \mathrm{Hom}_{\mathcal{C}_Q}(-, v)$ consists of all morphisms with codomain v.

We need to prove that F is a sheaf if and only if the above diagram is an equalizer for each vertex v. Let's proceed by showing both directions of the implication.

($\Rightarrow$) First, suppose that for each vertex v, given any family $\{s_f \in F(u)\}_{f:u\to v}$ satisfying $F(g)(s_f) = s_{f\circ g}$ for all composable pairs $f : u \to v$ and $g : w \to u$, there exists a unique $s \in F(v)$ such that $F(f)(s) = s_f$ for all $f : u \to v$. We need to show that the diagram is an equalizer.

To establish that $\epsilon : F(v) \to \prod_{f:u\to v} F(u)$ is the equalizer of p_1 and p_2, we must verify:

1. $p_1 \circ \epsilon = p_2 \circ \epsilon$ (the diagram commutes) 2. For any morphism $\phi : X \to \prod_{f:u\to v} F(u)$ such that $p_1 \circ \phi = p_2 \circ \phi$, there exists a unique morphism $\tilde{\phi} : X \to F(v)$ such that $\epsilon \circ \tilde{\phi} = \phi$ (universal property)

For commutativity (1), let $s \in F(v)$. Then:

$$(p_1 \circ \epsilon)(s) = p_1((F(f)(s))_{f\in R_v}) \tag{5}$$

$$= (F(f \circ g)(s))_{f,g} \tag{6}$$

$$= (F(g)(F(f)(s)))_{f,g} \quad \text{(by functoriality of } F\text{)} \tag{7}$$

$$= p_2((F(f)(s))_{f\in R_v}) \tag{8}$$

$$= (p_2 \circ \epsilon)(s) \tag{9}$$

For the universal property (2), consider any morphism $\phi : X \to \prod_{f:u\to v} F(u)$ such that $p_1 \circ \phi = p_2 \circ \phi$. For an element $x \in X$, we get a family $\phi(x) = (s_f)_{f\in R_v}$ where each $s_f \in F(u)$ for $f : u \to v$. The condition $p_1 \circ \phi = p_2 \circ \phi$ implies:

$$p_1((s_f)_{f\in R_v}) = p_2((s_f)_{f\in R_v}) \tag{10}$$

$$(s_{f\circ g})_{f,g} = (F(g)(s_f))_{f,g} \tag{11}$$

Therefore, for all composable pairs $f : u \to v$ and $g : w \to u$, we have $s_{f\circ g} = F(g)(s_f)$. By our assumption, there exists a unique $s \in F(v)$ such that $F(f)(s) = s_f$ for all $f : u \to v$. Define $\tilde{\phi}(x) = s$. This gives us a map $\tilde{\phi} : X \to F(v)$.

To verify $\epsilon \circ \tilde{\phi} = \phi$, we compute:

$$\begin{aligned}
(\epsilon \circ \tilde{\phi})(x) &= \epsilon(\tilde{\phi}(x)) && (12)\\
&= \epsilon(s) && (13)\\
&= (F(f)(s))_{f\in R_v} && (14)\\
&= (s_f)_{f\in R_v} && (15)\\
&= \phi(x) && (16)
\end{aligned}$$

For uniqueness of $\tilde{\phi}$, suppose $\tilde{\phi}' : X \to F(v)$ also satisfies $\epsilon \circ \tilde{\phi}' = \phi$. Then for any $x \in X$:

$$\begin{aligned}
(F(f)(\tilde{\phi}'(x)))_{f\in R_v} &= (\epsilon \circ \tilde{\phi}')(x) && (17)\\
&= \phi(x) && (18)\\
&= (\epsilon \circ \tilde{\phi})(x) && (19)\\
&= (F(f)(\tilde{\phi}(x)))_{f\in R_v} && (20)
\end{aligned}$$

In particular, for $f = \mathrm{id}_v$, we get $F(\mathrm{id}_v)(\tilde{\phi}'(x)) = F(\mathrm{id}_v)(\tilde{\phi}(x))$. Since $F(\mathrm{id}_v) = \mathrm{id}_{F(v)}$, we have $\tilde{\phi}'(x) = \tilde{\phi}(x)$ for all $x \in X$, thus $\tilde{\phi}' = \tilde{\phi}$.

($\Leftarrow$) Now, suppose the diagram is an equalizer for each vertex v. We need to show that for any family $\{s_f \in F(u)\}_{f:u\to v}$ satisfying $F(g)(s_f) = s_{f\circ g}$ for all composable pairs $f : u \to v$ and $g : w \to u$, there exists a unique $s \in F(v)$ such that $F(f)(s) = s_f$ for all $f : u \to v$.

Let $\{s_f \in F(u)\}_{f:u\to v}$ be such a family. Define $\alpha = (s_f)_{f\in R_v} \in \prod_{f:u\to v} F(u)$. The compatibility condition $F(g)(s_f) = s_{f\circ g}$ is precisely the assertion that $p_1(\alpha) = p_2(\alpha)$.

Since the diagram is an equalizer, there exists a unique $s \in F(v)$ such that $\epsilon(s) = \alpha$. By definition of ϵ, we have $\epsilon(s) = (F(f)(s))_{f\in R_v}$. Hence, $F(f)(s) = s_f$ for all $f : u \to v$, which is what we needed to prove.

Therefore, the simplified sheaf condition is equivalent to the standard sheaf condition for our specific topology J. $\square$

Example 3 (A non-sheaf)**.** *Consider a presheaf F on the path category of a graph with two vertices*

a, b and one edge $e : a \to b$. Suppose:

- *$F(a) = F(b) = k$*
- *$F(e) : F(b) \to F(a)$ is the zero map*
- *$F(\mathrm{id}_a) = F(\mathrm{id}_b) = \mathrm{id}_k$*

This F is not a sheaf. To see why, consider the covering sieve $R_b = \{\mathrm{id}_b, e\}$. Given sections $s_{\mathrm{id}_b} = 1 \in F(b)$ and $s_e = 0 \in F(a)$, we need to check compatibility:

$$F(\mathrm{id}_a)(s_e) = s_{e \circ \mathrm{id}_a} = s_e = 0$$

$$F(e)(s_{\mathrm{id}_b}) = 0 \neq 1 = s_{\mathrm{id}_b \circ e} = s_e$$

The compatibility fails, so this family cannot be glued.

2.2 Sheaves via dualization

We now show how to construct sheaves from covariant functors (quiver representations) via dualization.

Definition 4 (Dual presheaf)**.** *Given a covariant functor $V : \mathcal{C}_Q \to \mathbf{Vect}_k$, we define the* dual presheaf *$F = V^* : \mathcal{C}_Q^{op} \to \mathbf{Vect}_k$ by:*

- *$F(v) = V(v)^* = \mathrm{Hom}_k(V(v), k)$ for each vertex v*
- *$F(\alpha) = V(\alpha)^* : V(y)^* \to V(x)^*$ for each morphism $\alpha : x \to y$*

where $V(\alpha)^*$ is the dual (transpose) of the linear map $V(\alpha)$.

Remark 2. *Explicitly, for $\phi \in V(y)^*$ and $\alpha : x \to y$, we have:*

$$(F(\alpha))(\phi) = V(\alpha)^*(\phi) = \phi \circ V(\alpha) \in V(x)^*$$

This means $F(\alpha)$ pulls back linear functionals on $V(y)$ to linear functionals on $V(x)$.

Theorem 3. *Let $V : \mathcal{C}_Q \to \mathbf{Vect}_k$ be a functor and $F = V^*$ the dual presheaf. Then F is a sheaf on the site $(\mathcal{C}_Q, J)$.*

Proof. We need to verify the sheaf condition for each vertex v. Let $\{s_f \in F(u)\}_{f:u \to v}$ be a compatible family, meaning that for all composable $f : u \to v$ and $g : w \to u$:

$$F(g)(s_f) = s_{f \circ g}$$

Since $F(u) = V(u)^*$, each s_f is a linear functional $s_f : V(u) \to k$. The compatibility condition becomes:

$$s_f \circ V(g) = s_{f \circ g}$$

We need to find a unique $s \in F(v) = V(v)^*$ such that $F(f)(s) = s_f$ for all $f : u \to v$. This means we need $s \circ V(f) = s_f$.

Define $s : V(v) \to k$ as follows. For any $x \in V(v)$, since $R_v = \mathrm{Hom}(-, v)$ contains all morphisms into v, including id_v, we have:

$$s(x) = s_{\mathrm{id}_v}(x)$$

We need to verify this is well-defined and satisfies the required properties.

Well-definedness: This is immediate since s_{id_v} is given as part of the compatible family.

Verification: For any $f : u \to v$, we need to check that $s \circ V(f) = s_f$. For any $y \in V(u)$:

$$\begin{aligned}
(s \circ V(f))(y) &= s(V(f)(y)) && (21)\\
&= s_{\mathrm{id}_v}(V(f)(y)) && (22)\\
&= (s_{\mathrm{id}_v} \circ V(f))(y) && (23)\\
&= s_{\mathrm{id}_v \circ f}(y) \quad \text{(by compatibility)} && (24)\\
&= s_f(y) && (25)
\end{aligned}$$

Uniqueness: If $s' \in V(v)^*$ also satisfies $s' \circ V(f) = s_f$ for all f, then in particular for $f = \mathrm{id}_v$:

$$s' = s' \circ V(\mathrm{id}_v) = s_{\mathrm{id}_v} = s$$

Therefore, $F = V^*$ is a sheaf. □

Example 4 (Concrete computation)**.** *Let Q be the graph with vertices $\{a, b\}$ and one edge $e : a \to b$. Let V be the representation with $V(a) = V(b) = k$ and $V(e) = \mathrm{id}_k$. Then:*

- $F(a) = F(b) = k^* \cong k$
- $F(e) : F(b) \to F(a)$ *is* $(\mathrm{id}_k)^* = \mathrm{id}_{k^*} \cong \mathrm{id}_k$

To verify F is a sheaf, consider the covering sieve $R_b = \{\mathrm{id}_b, e\}$. A compatible family consists of $s_{\mathrm{id}_b} \in F(b)$ and $s_e \in F(a)$ such that:

$$F(\mathrm{id}_a)(s_e) = s_{e \circ \mathrm{id}_a} = s_e$$

$$F(e)(s_{\mathrm{id}_b}) = s_{\mathrm{id}_b} = s_{\mathrm{id}_b \circ e} = s_e$$

The second equation forces $s_{\mathrm{id}_b} = s_e$. Given such a family, the unique glued section is $s = s_{\mathrm{id}_b} \in F(b)$.

2.3 Properties of the sheaf construction

Theorem 4 (Functoriality)**.** *The assignment $V \mapsto V^*$ defines a contravariant functor from the category of representations* $\mathbf{Rep}(Q)$ *to the category of sheaves* $\mathbf{Sh}(\mathcal{C}_Q, J)$*.*

Proof. To establish that $V \mapsto V^*$ is a contravariant functor, we need to verify:

1. It maps objects in $\mathbf{Rep}(Q)$ to objects in $\mathbf{Sh}(\mathcal{C}_Q, J)$
2. It maps morphisms in $\mathbf{Rep}(Q)$ to morphisms in $\mathbf{Sh}(\mathcal{C}_Q, J)$ with direction reversed
3. It preserves identity morphisms
4. It preserves composition in the reversed order

For (1), we have already established in Theorem 4 that if V is a representation, then V^* is indeed a sheaf on $(\mathcal{C}_Q, J)$. For (2), let $\eta : V \to W$ be a morphism of representations. By definition, η is a natural transformation between functors $V, W : \mathcal{C}_Q \to \mathbf{Vect}_k$. This means that for every morphism $\alpha : x \to y$ in $\mathcal{C}_Q$, the following diagram commutes:

$$\begin{tikzcd} V(x) \arrow[r, "V(\alpha)"] \arrow[d, "\eta_x"'] & V(y) \arrow[d, "\eta_y"] \\ W(x) \arrow[r, "W(\alpha)"'] & W(y) \end{tikzcd}$$

We define $\eta^* : W^* \to V^*$ as follows: for each vertex $v \in Q_0$, let $\eta_v^* : W(v)^* \to V(v)^*$ be given by $\eta_v^*(\phi) = \phi \circ \eta_v$ for all $\phi \in W(v)^*$.

To show that η^* is a morphism of sheaves, we must verify that it is a natural transformation between the presheaves W^* and V^*. That is, for every morphism $\alpha : x \to y$ in $\mathcal{C}_Q$, the following diagram must commute:

$$\begin{tikzcd} W(y)^* \arrow[r, "W^*(\alpha)"] \arrow[d, "\eta_y^*"'] & W(x)^* \arrow[d, "\eta_x^*"] \\ V(y)^* \arrow[r, "V^*(\alpha)"'] & V(x)^* \end{tikzcd}$$

Let $\phi \in W(y)^*$. Following the upper-right path:

$$\begin{aligned} (\eta_x^* \circ W^*(\alpha))(\phi) &= \eta_x^*(W^*(\alpha)(\phi)) && (26) \\ &= \eta_x^*(\phi \circ W(\alpha)) && (27) \\ &= (\phi \circ W(\alpha)) \circ \eta_x && (28) \\ &= \phi \circ (W(\alpha) \circ \eta_x) && (29) \end{aligned}$$

Following the lower-left path:

$$\begin{aligned} (V^*(\alpha) \circ \eta_y^*)(\phi) &= V^*(\alpha)(\eta_y^*(\phi)) && (30) \\ &= V^*(\alpha)(\phi \circ \eta_y) && (31) \\ &= (\phi \circ \eta_y) \circ V(\alpha) && (32) \\ &= \phi \circ (\eta_y \circ V(\alpha)) && (33) \end{aligned}$$

By the naturality of η, we have $W(\alpha) \circ \eta_x = \eta_y \circ V(\alpha)$. Therefore,

$$\phi \circ (W(\alpha) \circ \eta_x) = \phi \circ (\eta_y \circ V(\alpha)) \quad (34)$$

which establishes that η^* is a natural transformation.

Since both V^* and W^* are sheaves (by Theorem 4), and η^* is a natural transformation between the underlying presheaves, it is a morphism in $\mathbf{Sh}(\mathcal{C}_Q, J)$.

For (3), we need to show that $(id_V)^* = id_{V^*}$. Let $id_V : V \to V$ be the identity morphism on V. For any vertex v and any $\phi \in V(v)^*$, we have:

$$\begin{aligned}
(id_V)_v^*(\phi) &= \phi \circ (id_V)_v && (35)\\
&= \phi \circ id_{V(v)} && (36)\\
&= \phi && (37)\\
&= id_{V(v)^*}(\phi) && (38)
\end{aligned}$$

Therefore, $(id_V)^* = id_{V^*}$.

For (4), let $\eta : V \to W$ and $\xi : W \to Z$ be morphisms in $\mathbf{Rep}(Q)$. We need to show that $(\xi \circ \eta)^* = \eta^* \circ \xi^*$. For any vertex v and any $\phi \in Z(v)^*$, we compute:

$$\begin{aligned}
(\xi \circ \eta)_v^*(\phi) &= \phi \circ (\xi \circ \eta)_v && (39)\\
&= \phi \circ (\xi_v \circ \eta_v) && (40)\\
&= (\phi \circ \xi_v) \circ \eta_v && (41)\\
&= \eta_v^*(\phi \circ \xi_v) && (42)\\
&= \eta_v^*(\xi_v^*(\phi)) && (43)\\
&= (\eta_v^* \circ \xi_v^*)(\phi) && (44)\\
&= (\eta^* \circ \xi^*)_v(\phi) && (45)
\end{aligned}$$

Thus, $(\xi \circ \eta)^* = \eta^* \circ \xi^*$, which confirms the contravariance of the functor.

Having verified conditions (1)-(4), we conclude that $V \mapsto V^*$ is indeed a contravariant functor from $\mathbf{Rep}(Q)$ to $\mathbf{Sh}(\mathcal{C}_Q, J)$. □

Remark 3 (Limitations of the construction)**.** *While we have established that dual presheaves are sheaves for our coarse topology, this construction essentially mirrors the original representation structure. The topology J is so coarse that the sheaf condition doesn't impose significant new constraints. This motivates considering finer topologies, as discussed in the conclusion.*

3 The Discrete Topology on $\mathcal{C}_Q$

3.1 Definition of the discrete topology

In contrast to the coarse topology discussed previously, we now consider the finest topology on the path category $\mathcal{C}_Q$, known as the *discrete topology.*

Definition 5 (Discrete topology)**.** *The discrete topology J_{disc} on $\mathcal{C}_Q$ is defined by declaring that every sieve is a covering sieve. That is, for every object $v \in Q_0$ and every sieve S on v, we have*

$S \in J_{disc}(v)$.

Remark 4. *The discrete topology is the finest possible topology on any category - it makes every sieve covering, which imposes maximal constraints on sheaves. This is distinct from the atomic topology (where only sieves containing the identity are covering).*

Before proceeding, we clarify what it means for one topology to be finer than another:

Definition 6 (Finer topology)**.** *Let J and J' be two topologies on a category $\mathcal{C}$. We say that J is* finer *than J' if and only if for every object U of $\mathcal{C}$ we have $J'(U) \subseteq J(U)$, which is to say that any covering sieve of J' is also a covering sieve of J.*

3.2 Verification of the Grothendieck topology axioms

Theorem 5. *The discrete topology J_{disc} satisfies the Grothendieck topology axioms on $\mathcal{C}_Q$.*

Proof. We verify axioms (GT1)–(GT3):
(GT1) Existence of maximal cover: The maximal sieve $\mathrm{Hom}_{\mathcal{C}_Q}(-, v)$ is declared covering for each vertex v, so (GT1) holds.
(GT2) Stability under pullback: Let S be any covering sieve on X and $f : Y \to X$ any morphism. In the discrete topology, every sieve is covering by definition. The pullback sieve is:

$$f^*(S) = \{g : Z \to Y \mid f \circ g \in S\}$$

Since every sieve is covering in J_{disc}, we automatically have $f^*(S) \in J_{\mathrm{disc}}(Y)$.
(GT3) Transitivity: Let S be a covering sieve on X and R any sieve on X such that for all $h : Y \to X$ in S, we have $h^*(R) \in J_{\mathrm{disc}}(Y)$. Since every sieve is covering in the discrete topology, we have $R \in J_{\mathrm{disc}}(X)$ by definition. □

3.3 Comparison with the coarse topology

Proposition 6 (Topology refinement)**.** *The discrete topology J_{disc} is finer than the coarse topology J defined previously. That is, for every vertex v and every sieve S:*

$$S \in J(v) \Rightarrow S \in J_{disc}(v)$$

Proof. In the coarse topology, the only covering sieve on vertex v is the maximal sieve $R_v = \mathrm{Hom}_{\mathcal{C}_Q}(-, v)$. Since the discrete topology declares every sieve to be covering, we certainly have $R_v \in J_{\mathrm{disc}}(v)$. More generally, since J_{disc} contains all sieves as covering sieves, it contains all covering sieves from J. □

4 Sheaves on the Discrete Site

4.1 Characterization of sheaves

The discrete topology has a dramatically different effect on the sheaf condition compared to the coarse topology.

Theorem 7 (Discrete sheaf condition)**.** *A presheaf $F : \mathcal{C}_Q^{\mathrm{op}} \to$ **Vect**$_k$ is a sheaf on $(\mathcal{C}_Q, J_{disc})$ if and only if for every vertex v and every morphism $f : u \to v$, the restriction map $F(f) : F(v) \to F(u)$ is an isomorphism.*

Proof. We prove both directions of the equivalence.

($\Leftarrow$) If all restriction maps are isomorphisms, then F is a sheaf:

Suppose $F(f) : F(v) \to F(u)$ is an isomorphism for every morphism $f : u \to v$ in $\mathcal{C}_Q$. We need to verify the sheaf condition for every covering sieve in J_{disc}.

Since we are in the discrete topology, every sieve is covering. Let S be any sieve on vertex v. The sheaf condition requires that the diagram

$$F(v) \xrightarrow{\epsilon} \prod_{(f:w\to v)\in S} F(w) \rightrightarrows \prod_{\substack{(f:w\to v)\in S \\ (g:z\to w)}} F(z)$$

is an equalizer, where:

- $\epsilon(s) = (F(f)(s))_{f\in S}$
- The two parallel arrows p_1, p_2 are defined by the compatibility conditions for families of sections

Since every sieve is covering in the discrete topology, in particular the sieve $\{\mathrm{id}_v\}$ is covering. For this minimal sieve, the sheaf condition becomes: the map

$$F(v) \xrightarrow{\epsilon} F(v)$$

given by $\epsilon(s) = F(\mathrm{id}_v)(s) = s$ is an equalizer of the empty family of parallel morphisms. This is automatically satisfied.

For any larger sieve S containing id_v, the sheaf condition follows from the fact that $F(\mathrm{id}_v) = \mathrm{id}_{F(v)}$ is an isomorphism, and all other restriction maps $F(f)$ for $f \in S$ are also isomorphisms by assumption. The equalizer condition is satisfied because we can uniquely reconstruct any section from its value at v via the isomorphisms $F(f)^{-1}$.

($\Rightarrow$) If F is a sheaf, then all restriction maps are isomorphisms:

Suppose F is a sheaf on $(\mathcal{C}_Q, J_{\text{disc}})$. We will show that every restriction map $F(f) : F(v) \to F(u)$ is an isomorphism by proving it is both injective and surjective.

Let $f : u \to v$ be any morphism in $\mathcal{C}_Q$. Consider the sieve $S = \{\mathrm{id}_v, f\}$ on vertex v. Since we are in the discrete topology, this sieve is covering.

The sheaf condition for this sieve states that the diagram

$$F(v) \xrightarrow{\epsilon} F(v) \times F(u) \rightrightarrows F(v) \times F(u) \times F(u)$$

is an equalizer, where:

- $\epsilon(s) = (s, F(f)(s))$ for $s \in F(v)$
- $p_1((s_{\mathrm{id}_v}, s_f)) = (s_{\mathrm{id}_v}, s_f, s_f)$ (using compatibility via id_u)
- $p_2((s_{\mathrm{id}_v}, s_f)) = (s_{\mathrm{id}_v}, s_f, F(\mathrm{id}_u)(s_f)) = (s_{\mathrm{id}_v}, s_f, s_f)$

Since $p_1 = p_2$, the equalizer condition simplifies to: ϵ is an isomorphism onto its image, and every element $(a, b) \in F(v) \times F(u)$ in the image of ϵ satisfies $b = F(f)(a)$.

Injectivity of $F(f)$: Suppose $F(f)(s_1) = F(f)(s_2)$ for $s_1, s_2 \in F(v)$. Then $\epsilon(s_1) = (s_1, F(f)(s_1))$ and $\epsilon(s_2) = (s_2, F(f)(s_2))$. Since the second components are equal and ϵ must be injective for the equalizer property, we need $s_1 = s_2$. Thus $F(f)$ is injective.

Surjectivity of $F(f)$: Let $t \in F(u)$ be arbitrary. Consider the family (s_{id_v}, s_f) where we want to find $s_{\mathrm{id}_v} \in F(v)$ such that $s_f = t$ and the compatibility conditions hold. The sheaf condition guarantees that such an s_{id_v} exists and is unique, and we must have $F(f)(s_{\mathrm{id}_v}) = s_f = t$. Therefore $F(f)$ is surjective.

Since $F(f)$ is both injective and surjective, it is an isomorphism.

This argument applies to every morphism $f : u \to v$ in $\mathcal{C}_Q$, completing the proof. □

Corollary 8. *A presheaf F is a sheaf on $(\mathcal{C}_Q, J_{disc})$ if and only if F is locally constant in the sense that all restriction maps are isomorphisms.*

4.2 Examples of discrete sheaves

Example 5 (Constant sheaf)**.** *For any vector space $V \in \boldsymbol{Vect}_k$, define the constant presheaf $\underline{V}$ by:*

- *$\underline{V}(v) = V$ for all vertices v*
- *$\underline{V}(\alpha) = id_V$ for all morphisms α*

This is automatically a sheaf on $(\mathcal{C}_Q, J_{disc})$ since all restriction maps are isomorphisms.

Proof. Every restriction map $\underline{V}(\alpha) = \mathrm{id}_V$ is clearly an isomorphism, so by the discrete sheaf condition, $\underline{V}$ is a sheaf on the discrete site. □

Example 6 (Locally constant sheaf)**.** *Let Q be a graph and suppose each connected component of Q has a distinguished "base vertex." For each connected component C, choose a vector space V_C. Define F by:*

- *$F(v) = V_C$ where C is the connected component containing v*

- *$F(\alpha) = id_{V_C}$ for any morphism α within component C*

This gives a sheaf that is constant on each connected component but may vary between components.

Proof. Since each restriction map is id_{V_C} for some component C, all restriction maps are isomorphisms, making F a sheaf by the discrete sheaf condition. $\square$

4.3 Non-examples

Example 7 (Non-sheaf in discrete topology)**.** *Consider the graph with vertices $\{a, b\}$ and edge $e : a \to b$. Define a presheaf F by:*

- *$F(a) = k$, $F(b) = k^2$*
- *$F(e) : k^2 \to k$ given by $(x, y) \mapsto x$*

This is not a sheaf on $(\mathcal{C}_Q, J_{disc})$ because $F(e)$ is not an isomorphism.

Proof. The map $F(e) : k^2 \to k$ given by $(x, y) \mapsto x$ is not injective (since $(1, 0)$ and $(1, 1)$ both map to 1), hence not an isomorphism. By the discrete sheaf condition, F cannot be a sheaf on the discrete site. $\square$

5 Comparison Between Coarse and Discrete Topologies

5.1 Sheaf categories

Let $\mathbf{Sh}(\mathcal{C}_Q, J)$ denote the category of sheaves on the coarse site and $\mathbf{Sh}(\mathcal{C}_Q, J_{\mathrm{disc}})$ the category of sheaves on the discrete site.

Theorem 9 (Inclusion of sheaf categories)**.** *There is a natural inclusion functor:*

$$\iota : \boldsymbol{Sh}(\mathcal{C}_Q, J_{disc}) \to \boldsymbol{Sh}(\mathcal{C}_Q, J)$$

Moreover, this functor is fully faithful.

Proof. Since J_{disc} is the finest possible topology and J is coarser, we have $J(v) \subseteq J_{\mathrm{disc}}(v)$ for every object v. Therefore, every sheaf for the discrete topology automatically satisfies the sheaf condition for the coarse topology: if F satisfies the sheaf condition for all sieves (since all are covering in J_{disc}), then in particular it satisfies the sheaf condition for the covering sieves in J.

The functor ι is the inclusion of the full subcategory of locally constant sheaves within all sheaves on the coarse site.

Full faithfulness follows because morphisms between sheaves are natural transformations between the underlying presheaves, and these are the same regardless of which topology we consider. That is, if F, G are discrete sheaves, then

$$\mathrm{Hom}_{\mathbf{Sh}(\mathcal{C}_Q, J_{\mathrm{disc}})}(F, G) = \mathrm{Hom}_{\mathbf{Sh}(\mathcal{C}_Q, J)}(\iota(F), \iota(G))$$

since the morphisms are determined by the underlying presheaf structure. $\square$

5.2 Relationship to quiver representations

Proposition 10. *The category* $\textbf{\textit{Sh}}(\mathcal{C}_Q, J_{disc})$ *is equivalent to the category of representations of the groupoid completion of* Q*.*

Proof. A sheaf on $(\mathcal{C}_Q, J_{\text{disc}})$ assigns isomorphisms to all morphisms in $\mathcal{C}_Q$ (by the discrete sheaf condition). This is equivalent to giving a representation of the groupoid obtained by formally inverting all morphisms in $\mathcal{C}_Q$.

More precisely, let $\mathcal{G}_Q$ denote the groupoid completion of $\mathcal{C}_Q$, which has the same objects as $\mathcal{C}_Q$ but where every morphism has been made invertible. A functor $F : \mathcal{G}_Q \to \mathbf{Vect}_k$ corresponds to a contravariant functor $F : \mathcal{C}_Q^{\text{op}} \to \mathbf{Vect}_k$ where all restriction maps are isomorphisms, which is precisely a discrete sheaf. $\square$

5.3 Functorial relationships

Theorem 11 (Restriction and extension functors)**.** *There exist adjoint functors:*

$$\textbf{\textit{Sh}}(\mathcal{C}_Q, J_{disc}) \underset{R}{\overset{L}{\rightleftarrows}} \textbf{\textit{Sh}}(\mathcal{C}_Q, J)$$

where L *is left adjoint to the restriction functor* $R = \iota$*.*

Proof. The restriction functor R is the inclusion ι. Its left adjoint L can be constructed as follows: given a sheaf F on the coarse site, define LF by:

$$(LF)(v) = \varinjlim_{f:u \to v} F(u)$$

where the colimit is taken over all morphisms into v. The universal property of colimits ensures that LF satisfies the discrete sheaf condition.

To verify this is an adjunction, we need to show there is a natural isomorphism:

$$\text{Hom}_{\mathbf{Sh}(\mathcal{C}_Q, J_{\text{disc}})}(LF, G) \cong \text{Hom}_{\mathbf{Sh}(\mathcal{C}_Q, J)}(F, RG)$$

Given a morphism $\phi : LF \to G$ in $\mathbf{Sh}(\mathcal{C}_Q, J_{\text{disc}})$, we can construct $\tilde{\phi} : F \to RG$ using the universal property of the colimit defining $LF(v)$. Conversely, given $\psi : F \to RG$, we can construct $\hat{\psi} : LF \to G$ by the universal property. These constructions are inverse and natural in F and G. $\square$

5.4 Ramifications for representation theory

The discrete topology provides a different perspective on quiver representations:

- **Coarse topology:** Sheaves correspond to dual representations via dualization
- **Discrete topology:** Sheaves correspond to locally constant assignments of vector spaces

This suggests that intermediate topologies could interpolate between these extremes, potentially capturing different aspects of representation-theoretic structure.

6 Future Research Questions and Intermediate Topologies

6.1 Edge-generated topology

An intermediate topology between the coarse and discrete cases could be defined by declaring that a sieve S on vertex v is covering if and only if it contains at least one morphism that factors through an edge terminating at v.

Definition 7 (Edge-generated topology)**.** *A sieve S on vertex v is covering in the edge-generated topology J_{edge} if there exists an edge $e \in Q_1$ with $t(e) = v$ such that $e \in S$.*

This topology would distinguish between vertices based on their incoming edge structure and could provide a more refined analysis of quiver representations. It sits properly between the coarse and discrete topologies: it is finer than the coarse topology (since it requires actual edges, not just any morphism) but coarser than the discrete topology (since not every sieve is covering).

6.2 Graded topologies

For applications to graded representations, one could define topologies that respect path length:

Definition 8 (Length-graded topology)**.** *A sieve S on vertex v is covering in the n-graded topology if it contains all morphisms of length at most n with codomain v.*

This could be useful for studying filtered or graded module categories associated to quivers.

7 Conclusion

We have constructed two fundamental Grothendieck topologies on path categories of finite graphs, establishing connections between sheaf theory and quiver representation theory. The coarse topology J provides a natural setting where quiver representations dualize to sheaves, while the discrete topology J_{disc} forces sheaves to be locally constant, corresponding to groupoid representations.

These topologies represent extremes with markedly different behaviors. Coarse sheaves arise systematically from representations through dualization, preserving structural information while enabling sheaf-theoretic techniques. Discrete sheaves, being locally constant, capture rigid aspects where all restriction maps are isomorphisms. The functorial relationships between these categories - particularly the inclusion functor and its left adjoint - demonstrate these form part of a coherent framework rather than isolated constructions.

Our work points toward several research directions. Intermediate topologies between coarse and discrete extremes could capture subtler representation-theoretic phenomena. Edge-generated topologies might distinguish morphisms factoring through graph edges, while graded topologies could respect path length. Such constructions may provide tools for studying stability conditions, persistent homology over quiver categories, or derived sheaf categories.

The sheaf-theoretic perspective offers computational advantages for the coarse topology and conceptual insights for the discrete case. This framework may prove valuable as representation theory finds new applications, providing different analytical tools while potentially revealing hidden structures in classical problems.

By placing quiver representations within Grothendieck's general framework, we open access to topos-theoretic methods while suggesting that topology choice significantly influences algebraic properties. This work demonstrates how categorical thinking can illuminate classical objects, potentially leading to new classification results or structural relationships in representation theory.

The connection to representation theory is fundamental: quiver representations naturally arise in the study of finite-dimensional algebras, where they encode the module structure. Our sheaf-theoretic approach provides new tools for analyzing these representations, particularly through the dualization functor in the coarse case and the rigidity constraints in the discrete case. Future work could explore how these topological perspectives inform classical problems like Gabriel's theorem on representation-finite quivers or the study of derived categories of quiver representations.

References

[1] Saunders Mac Lane and Ieke Moerdijk. *Sheaves in Geometry and Logic: A First Introduction to Topos Theory*. Springer-Verlag, New York, 1992.

[2] Peter Gabriel. "Unzerlegbare Darstellungen I." *Manuscripta Mathematica* 6.1 (1972): 71–103.

[3] The Stacks Project Authors. *Stacks Project*. `https://stacks.math.columbia.edu`, 2023.

[4] Peter Webb. *Math 8300 Representations of Finite Dimensional Algebras*. University of Minnesota, 2019.