Quiver Signal Processing

Rayna Andreeva

University of Edinburgh

07/12/2020

Introduction

- Basics for a signal processing framework on quiver representations
- Can handle heterogeneous multidimensional information in networks.
- Provides a generalisation of graph neural networks (GNN)
- Useful when studying directed graphs

Introduction

- Graphs are limited for modeling systems whose agents not only pass messages in a heterogenous environment, but also process heterogenous data types
- Cellular sheaves are a data structure for stitching together assignments of data to both nodes and edges of a graph...
- They require "hidden" data assigned to edges, which could be inefficient or even unnatural as data must be pulled back from all edges incident to a node before it may be processed by a node

Introduction

- Quiver representation theory has a rich history in mathematics
- It has not been used much in engineering/computer science
- A quiver representation is an assignment of a vector space to each node and an assignment of a linear map associated to each arrow.
- A quiver representation provides a natural environment to handle heterogeneous information as vector spaces associated to nodes can be non-isomorphic i.e. have different dimensions

Quiver Representations

Definition 1
A quiver Q is a quadruple $Q = (Q_0, Q_1, h, t)$, consisting of sets Q_0 and Q_1, and maps $h, t : Q_1 \to Q_0$.

Some clarification: Q_0 are the nodes of Q and Q_1 are the arrows. For an arrow $a \in Q_1$, we say $h(a)$ is the head of a and $t(a)$ is the tail of a.

Quivers examples

Example 1

$$1 \xrightarrow{a_{1,2}} 2$$

Let $Q_0 = \{1, 2\}$, and $Q_1 = \{a_{1,2}\}$. The maps h and t are given by $h(a_{i,j}) = j$ and $t(a_{i,j}) = i$.

Example 2

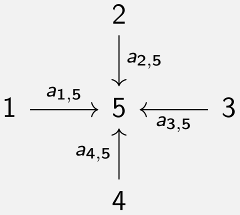

$Q_0 = \{1, 2, 3, 4, 5\}$, and $Q_1 = \{a_{1,5}, a_{2,5}, a_{3,5}, a_{4,5}\}$. The maps h and t are given by $h(a_{i,j}) = j$ and $t(a_{i,j}) = i$.

Quivers alone do not offer anything new to the GSP literature. However, the notion of a representation of a quiver provides the needed formal components used to associate information to a quiver in different ways.

Definition 2

A (finite-dimensional) representation π of a quiver $Q = (Q_0, Q_1, h, t)$ is an assignment of (finite-dimensional) vector spaces to nodes: $\pi : i \in Q_0 \to V_i$, and an assignment of linear maps to arrows: $\pi :$ $a \in Q_1 \to (\phi_{t(a),h(a)} : V_{t(a)} \to V_{h(a)})$.

Example (robotics)

- Quiver representations offer a formal tool to analyze and understand heterogeneous large-scale distributed robotic systems
- We can define an arbitrary graph where each node is a robot in the system
- To each node in the graph, we associate a vector space that decomposes as $V_{sense} \oplus V_{task}$ to represent the type of information that each robot can sense V_{sense}, and the configuration space of the task to be performed by the robot V_{task}.
- Message-passing between robots is given by linear operators that synthesize both sensing and task-related observations
- This setup makes sense when dividing robots in several species (where each species (robot type) is defined by the traits (capabilities) that it owns).

Reference - "Formalizing the impact of diversity on performance in a heterogeneous swarm of robots," in 2016 IEEE International Conference on Robotics and Automation (ICRA), 2016

When are two representations the same?

We start with a notion that will allow us to compare representations of a quiver.

Definition 3

Let π and ρ be representations of a quiver Q. An intertwining map $T : \pi \to \rho$ is a family of linear transformations $\{T_i : \pi(i) \to \rho(i) \mid i \in Q_0\}$ such that for every arrow $a \in Q_1$, the following diagram commutes:

$$\begin{array}{ccc} \pi(t(a)) & \xrightarrow{\pi(a)} & \pi(h(a)) \\ \downarrow{\scriptstyle T(t(a))} & & \downarrow{\scriptstyle T(h(a))} \\ \rho(t(a)) & \xrightarrow{\rho(a)} & \rho(h(a)). \end{array}$$

Example of isomorphic representations

Example: Consider the representations:

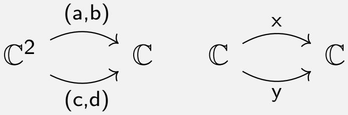

where $a, b, c, d, x, y \in \mathbb{C}$. If $(a, b) = (2, 1)$, $(c, d) = (6, 3)$, $x = 1$ and $y = 3$, construct a non-zero homomorphism $f : M \to N$.

Example of isomorphic representations, solution

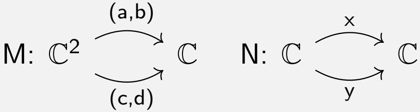

where $a, b, c, d, x, y \in \mathbb{C}$. If $(a, b) = (2, 1)$, $(c, d) = (6, 3)$, $x = 1$ and $y = 3$.

Since we have 2 nodes, we should construct 2 linear maps $f : \mathbb{C}^2 \to \mathbb{C}$ and $g : \mathbb{C} \to \mathbb{C}$ in such a way that both diagrams commute (we have 2 diagrams because of the 2 arrows):

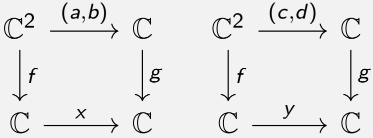

Example of isomorphic representations, solution

We define f and g by $f = (2, 1) : \mathbb{C}^2 \to \mathbb{C}$ and $g = 1 : \mathbb{C} \to \mathbb{C}$. Then

$$x \circ f = g \circ (a, b), y \circ f = g \circ (c, d).$$

Path in a quiver

Definition 4

A path γ in a quiver $Q = (Q_0, Q_1, h, t)$ of length $l \geq 1$ is a sequence $\gamma = a_l a_{l-1} \cdots a_1$ such that $a_i \in Q_1$ for all i and $t(a_{i+1}) = h(a_i)$ for $i = 1, \cdots l-1$ with $h(\gamma) = h(a_l)$ and $t(\gamma) = t(a_1)$. By convention, there is a trivial path e_i of length zero which can be defined by $h(e_i) = t(e_i) = i$ for every $i \in Q_0$. Let $Path(Q)$ denote the set of paths of a quiver Q.

Path and Path Algebra - but first, what is a field?

Definition 5
A field k is a set where sum and multiplication are defined.

- Define numbers and the operations we perform on them $\Longrightarrow$ Reals $k = \mathbb{R}$, complex numbers: $k = \mathbb{C}$
- We can define 2 operations, the sum and the product - you can add and multiply any two real numbers or any two complex numbers

Path and Path Algebra

Definition 6

Let k be a field (most common examples are $\mathbb{C}$ and $\mathbb{R}$). The path algebra of a quiver $Q = (Q_0, Q_1, h, t)$, denoted kQ, is the free vector space (meaning it has a basis) with basis $Path(Q)$ and a product on basis elements,

$$\gamma_2 . \gamma_1 = \begin{cases} \gamma_1 \gamma_2 & \mathrm{t}(\gamma_2) = h(\gamma_1) \\ 0 & \mathrm{t}(\gamma_2) \neq h(\gamma_1) \end{cases} \tag{1}$$

This definition provides a multiplicative operation for the algebra, which is concatenation when it makes sense - the first path's head matches the tail of the second path.

Associative Algebra

Definition 7

An associative algebra A is a vector space with a bilinear map $A \times A \to A$ mapping $(a, b) \to a * b$ and such that $(a * b) * c = a * (b * c)$

Algebra adds one more operation to a vector space - the possibility to multiply two vectors and stay within the vector space.

Representation of an algebra

Definition 8

A representation of an algebra A is given by a pair $(\mathcal{M}, \rho)$, where $\mathcal{M}$ is a vector space and ρ is a so-called homomorphism $\rho : A \to End(\mathcal{M})$, where $End(\mathcal{M})$ is the algebra of linear maps from a vector space $\mathcal{M}$ to itself. Notice that ρ is a homomorphism if ρ is a linear map, ρ respects products $\rho(ab) = \rho(a)\rho(b)$, and $\rho(1) = I$.

Note: when $\mathcal{M} = \mathbb{R}^n$, think of $End(\mathcal{M})$ as all the linear transformations, or the $n \times n$ square matrices.

Bijections between quiver representations and path algebras

Theorem 9

There is a bijection

$$\{\textit{Representations of } Q\} \leftrightarrow \{\textit{Representations of } kQ.\} \quad (2)$$

Explanation

One direction of the theorem implies that a given quiver representation π can be transformed into an equivalent representation of kQ, which consists of the data of a vector space $\mathcal{M}$ and a homomorphism ρ. Indeed $\mathcal{M} = \oplus_i V_i$ where $V_i = \pi(i)$ and $i \in Q_0$. We also point out that because the path algebra kQ is generated by the elements in $Path(Q)$, we can describe the action of ρ in terms of $\rho(p)$ for all $p \in Path(Q)$. In particular, the action of an element $p \in Path(Q)$ on an element $x \in \mathcal{M}$ is given by $y = \rho(p)x$ where

$$y(j) = \begin{cases} \pi(p)x(i) & if\,\mathrm{t(p) = i,\ h(p) = j} \\ 0 & otherwise, \end{cases} \tag{3}$$

where $\pi(p)$ for the path $p = a_l a_{l-1} \cdots a_1$ is given naturally by the composition map $\pi(p) = \pi(a_l)\pi(a_{l-1}). \cdots \pi(a_1)$.

Signal processing on quiver representations

We now proceed to develop a theory of signal processing on quiver representations under the umbrella of algebraic signal processing.

Definition 10

Let $Q = (Q_0, Q_1, h, t)$ be a quiver and π a representation of Q. Then a signal x on π is an element $x \in \oplus_{i \in Q_0} \pi(i)$. Additionally, the elements in kQ are called algebraic filters while their images in $End(\mathcal{M})$ via the homomorphism ρ associated to the representation of kQ are called quiver filters.

Then the filtering of a signal x is given by $\rho(c)x$ where $c \in kQ$. Notice that this is the convolution between the signal x and $\rho(c)$.

Example

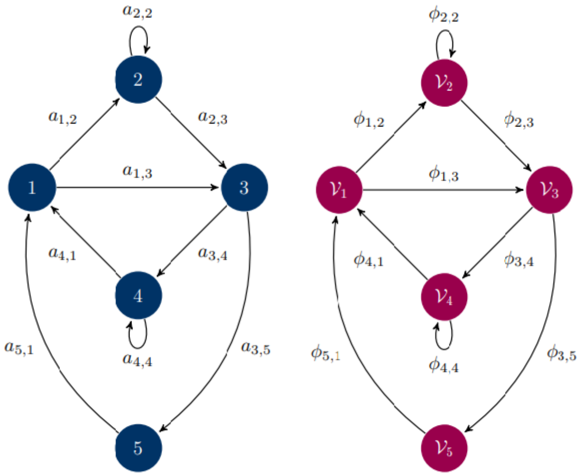

Figure 1: Quiver representation

Example

Taking into account the quiver depicted in on the left and its representation on the right, we have an example of filtering considering the algebraic filter $c = a_{5,1}a_{3,5} + a_{1,2}a_{4,1}a_{3,4} + a_{1,3}a_{2,3}$. Then, when filtering a signal $x \in \oplus_{i\in Q_0} V_i$, we have

$$y = \rho(c)x = \rho(a_{5,1}a_{3,5} + a_{1,2}a_{4,1}a_{3,4} + a_{1,3}a_{2,3})x. \qquad (4)$$

Using the linearity of the homomorphism ρ, we get:

$$y = \rho(a_{5,1}a_{3,5})x + \rho(a_{1,2}a_{4,1}a_{3,4})x + \rho(a_{1,3}a_{2,3})x. \qquad (5)$$

Then using equation (3) for the action of ρ on the elements of $\mathcal{M}$, we have $y(1) = \phi_{5,1}\phi_{3,5}x(3)$, $y(2) = \phi_{1,2}\phi_{4,1}\phi_{3,4}x(3)$ and $y(l) = 0$ for $l = 3, 4, 5$.

Shift operator

Definition 11

Let $Q = (Q_0, Q_1, h, t)$ be a quiver and a representation of Q. Let ρ be the homomorphism induced by π in a representation of the path algebra kQ. Then, the operator $\rho(p)$ for any $p \in Path(Q)$ is called a shift operator.

Indecomposable representation

Definition 12

A nontrivial representation π of a quiver Q is decomposable if π is isomorphic to $(\pi_1 \oplus \pi_2)$ for some nontrivial representations π_1, π_2 of Q. A representation that is not decomposable is called indecomposable.

Basic decomposition

Definition 13

Let $Q = (Q_0, Q_1, h, t)$ be a quiver, and consider θ_i, the indecomposable representations of Q. Then, we say that *pi* has a basic decomposition if

$$\pi \cong r_1\theta_1 \oplus r_2\theta_2 \cdots \oplus r_k\theta_k, \tag{6}$$

where $r_i\theta_i$ represents the direct sum of r_i copies of θ_i.

Example for basic decomposition

Let us consider the quiver $Q : \circ \longrightarrow \circ$, which consists only of two nodes and one arrow, and consider the representation π of Q given by $\mathbb{C}^{r+s} \xrightarrow{\phi_{1,2}} \mathbb{C}^{r+t}$. The three indecomposable representations of Q are:

- (Rep 1) $\mathbb{C}\xrightarrow{1}\mathbb{C}$
- (Rep 2) $0\xrightarrow{0}\mathbb{C}$
- (Rep 3) $\mathbb{C}\xrightarrow{0}0$

We can see that π can expressed as a direct sum of r copies of (Rep1), s copies of (Rep3) and t copies of (Rep2).

Conclusion

- We have laid the groundwork for a signal processing framework on quiver representations
- We introduced the notions of signals, filters and basic decompositions
- We went over a few basic examples of quiver representations
- An application area of robotics was suggested

Conclusion and thoughts?

- This signal processing framework provides a new tool for handling heterogenous data distributed across networks
- It can be used for new neural convolutional architectures - what are the next steps for Quiver Neural Networks (QNN)?
- Does it have an advantage over GNNs?
- What are other application areas?
- High level of abstraction (but great for generalisation) - do theories like this one have a place in ML?
- Would this be computationally feasible?

Thank you!

Thank you for your attention!