

Set-to-Sequence Methods in Machine Learning: a Review

Mateusz Jurewicz

*Department of Computer Science
IT University of Copenhagen, 2300 Copenhagen, Denmark
Tjek A/S, 1408 Copenhagen, Denmark*

MAJU@ITU.DK, MJ@TJEK.COM

Leon Derczynski

*Department of Computer Science
IT University of Copenhagen, 2300 Copenhagen, Denmark*

LEOD@ITU.DK

Abstract

Machine learning on sets towards sequential output is an important and ubiquitous task, with applications ranging from language modelling and meta-learning to multi-agent strategy games and power grid optimization. Combining elements of representation learning and structured prediction, its two primary challenges include obtaining a meaningful, permutation invariant set representation and subsequently utilizing this representation to output a complex target permutation. This paper provides a comprehensive introduction to the field as well as an overview of important machine learning methods tackling both of these key challenges, with a detailed qualitative comparison of selected model architectures.

1. Introduction

We begin by providing a definition of the set-to-sequence field and outline its importance in various areas of application.

1.1 What is Set-to-Sequence?

Set-to-sequence encompasses a group of problems where input takes the form of unordered collections of elements and the output is an ordered sequence. These challenges can be approached as a machine learning problem, where models learn arbitrary functions for performing the set-to-sequence mapping.

Set-to-sequence covers combinatorial optimization and structure prediction problems where exhaustive search is often not tractable. Machine learning (ML) approaches to set-to-sequence combine set-encoding techniques with permutation learning and have found an exceptionally wide range of practical applications.

Many of the successful deep learning approaches take advantage of the structure in their input data. However, sets do not possess the kind of internal structure that images and natural language sentences do. In set-to-sequence our input data does not have an inherent order and therefore our models must take into consideration the *permutation invariance* of sets. Obtaining meaningful permutation invariant representations is an important challenge for machine learning models in order to ensure that the same set will not result in two different outputs, due to the arbitrary initial order in which its elements were presented to the model.

1.2 Why Does Set-to-Sequence Matter?

Machine learning set-to-sequence methods can approximate solutions to computationally intractable problems in many areas. They have been applied to learning competitive solvers for the NP-Hard Travelling Salesman Problem (Vinyals et al., 2015); tackling prominent NLP challenges such as

sentence ordering (Wang & Wan, 2019) and text summarization (Sun et al., 2019); and in multi-agent reinforcement learning (Sunehag et al., 2018). A notable example is the agent employed by the AlphaStar model, which defeated a grandmaster level player in the strategy game of Starcraft II (Vinyals et al., 2019).

Set-to-sequence ML models also play an important role in data-intensive 3D point cloud processing (Qi et al., 2017) and meta-learning (Huang et al., 2018b). Set-input and set-ordering problems themselves are prominent in a wide array of applications ranging from power grid optimization (Cui et al., 2019), where solving them led to power usage savings of up to 30%, through anomaly detection (Jung et al., 2015) to measurements of contaminated galaxy clusters (Ntampaka et al., 2016).

This review contributes to the field in two primary ways:

1. By providing a single point of entry for researchers interested in the set-to-sequence field and for applied practitioners solving set-input ordering challenges.
2. By comparing all the discussed methods via a number of aspects relevant for both academic and applied work and presenting this comparison in the form of easy-to-read tables, which will help guide the reader towards the most applicable method for their specific area of interest.

The remainder of this paper is structured in the following way: firstly, we introduce the reader to the necessary background concepts and related work in section 2, including specific notes on the adopted notation. Secondly, an overview of set encoding methods is given in section 3, with comparison tables and details of the underlying mathematical transformations. Thirdly, section 4 contains a survey of popular ordering methods, which use the encoded set representation to output a complex permutation. The lists provided in sections 3 and 4 are not exhaustive and focus primarily on deep learning approaches. Finally, a discussion of current limitations and directions for further research is given in section 5, followed by a short conclusive summary in section 6.

2. Background

In this section the reader is introduced to the key concepts related to machine learning on sets and permutation learning, with minor notes on notation throughout the rest of the paper. Additionally, a comprehensive overview of related work from other fields of machine learning is given, including natural language processing, information retrieval and set segmentation.

2.1 Important Concepts

For the purposes of this review a *set* can be intuitively defined as a collection of distinct elements, without a canonical order between them (Halmos, 2017). An important property of sets is that they can have a varying number of elements, also referred to as their *cardinality*. Whilst the intuitive definition of a set is susceptible to known paradoxes (Rang & Thomas, 1981), the machine learning methods discussed here do not require an in-depth understanding of the proper definition from axiomatic set theory. Interested readers can find further information pertaining to it in other referred publications (Takeuti & Zaring, 2013).

As per the axiom of extensionality, sets are defined only by their elements (Hayden et al., 1968). In practice this means that given, for example a set $A = \{x, y, z\}$ and set $B = \{z, y, x\}$ we know that $A = B$. The order in which the elements are presented in roster notation does not matter. From now on, when we refer to a set, we specifically limit our considerations to finite sets only.

Set-to-sequence ML methods are distinctly different from earlier, encoder-decoder sequence-to-sequence model architectures (Sutskever et al., 2014). The difference stems from having to handle set-input data. This imposes two requirements on set-to-sequence ML methods that are not met by most neural network models:

1. *Permutation Invariance*

The output of the model must be the same under every possible permutation of the elements from the input set.

2. *Varying Input Length*

The same model must be able to process input sets of different lengths.

If these criteria are not met, the ML model by definition treats its input as a sequence, not a set. Fully feed-forward methods fail to meet both criteria and the recurrent neural networks (RNNs), which form the foundation of most sequence-to-sequence autoencoders, are sensitive to alterations of the order of their input (Vinyals et al., 2016). To truly treat input data as an inherently unordered set we must be certain that permuting the input will not result in a different encoded set representation (Zaheer et al., 2017). Additionally, depending on the presence and type of a downstream task that uses this representation, we are interested in whether the final output is also permutation invariant, which is not necessarily the case with all reordering methods.

More formally, a function $f : \mathcal{P}(X) \rightarrow Y$ is permutation invariant regarding the order of the elements of its input set if for every permutation π the following property holds: $f(\{x_1, \dots, x_n\}) = f(\{x_{\pi(1)}, \dots, x_{\pi(n)}\})$. A related property of functions on sets, which has been formally investigated by Zaheer et al. (2017), is *permutation equivariance*. In tasks where each set element has an associated target label, such that these individual labels depend on the entirety of the set, we would ideally want our predicted labels to remain the same per element, regardless of how the original input set is permuted. That property is permutation equivariance.

At this point it is important to distinguish between two types of set-to-sequence challenges. In the first type the output is a reordering of the input elements, with the possibility of repeating an element multiple times in the output sequence or skipping it entirely. Such permutations with potential repetition and exclusion are further denoted as *complex permutations*. We can refer to the type of problems involving various kinds of permutations of the input elements as *self-referential set-to-sequence challenges*.

The self-referential set-to-sequence domain includes classic combinatorial optimization problems and forms the majority of this review. There are many different ways to frame this task and formalize the resulting output, which are discussed in section 4. They include primarily pointer-based attention (4.2), the generation of permutation matrices (4.3) and ranking scores (4.4).

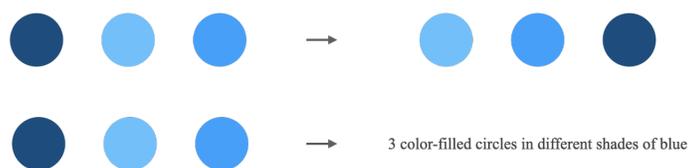


Figure 1: Set-to-Sequence Tasks by Referentiality

In the second type of set-to-sequence challenges, the task is to generate output that is sequential in nature, but is not defined as a permutation of the original input elements. We denote this as *non-referential set-to-sequence*. It encompasses for example summarization of a set of documents. Here, the input is indeed a set, with unique elements that do not have a canonical ordering to them. The output is a sequence of natural language tokens in the form of a human-readable summary, without referring directly to the elements of the input set. This area is only partially covered by this review.

For a simple visual explanation highlighting the difference between self-referential (top) and non-referential (bottom) set-to-sequence tasks, see Figure 1. In both cases the input is a set of 3 disks of varying shades of blue. In a self-referential setting the target might be a permutation from lightest to darkest. In a non-referential setting the target may be a sequence of word tokens describing the input set.

Specifically, all descriptions of set-input encoding methods are shared between the two types of set-to-sequence problems and therefore will be of value to readers interested in either. However, for considerations related to sequence prediction in general, areas of interest include recurrent neural network encoder-decoder models (Sutskever et al., 2014), reinforcement learning actor-critic methods (Bahdanau et al., 2017) and fully-connected transformer architectures as investigated by Vaswani et al. (2017), Devlin et al. (2019) and Brown et al. (2020).

In summary, the former type of set-to-sequence ML architectures, which this review focuses on, tackle two primary challenges:

1. Handling varying-length set-input data in a way that ensures permutation invariance
2. Handling outputs as complex permutations or *reorderings* of the original input

The first challenge, once solved, allows us to use machine learning methods to perform set-input regression, classification, recommendation (Vartak & Thiagarajan, 2017), as well as clustering problems and more (Edwards & Storkey, 2017).

Depending on the specific task at hand, the permutation invariant representation of the input set may also be required to encode higher order interactions between the input set elements, as seen in the work of Lee et al. (2019) and Zhang et al. (2020), which is a separate but important consideration in the area of encoding sets.

The second challenge of *permutation learning* is made simpler by solving the first one, but has also been successfully tackled without addressing it (Vinyals et al., 2015). It focuses on learning the proper order of arbitrary input elements. As a result, the model learns to predict the best structure of the output sequence composition.

2.2 Difficulty of Learning to Reorder

Permutation learning is an inherently difficult challenge. Even a relatively simple application of set-to-sequence methods to the Travelling Salesman Problem (TSP) in a two dimensional Euclidean space involves tackling an NP-Hard problem (Vinyals et al., 2015).

Whilst highly successful, polynomial time algorithms for obtaining good approximate solutions to this task do exist, such as the ones proposed by Arora (1996) and Karlin et al. (2020), it is also valuable to investigate the capacity of current machine learning techniques to learn them iteratively.

In the two dimensional Euclidean version of the TSP our input is a set of point coordinates and our desired output is a permutation of these points in a way that results in the shortest distance travelled between them. Additionally, we must not skip nor revisit any of the points. For a visual example, see Figure 2.

This challenge is difficult because the number of possible permutations increases factorially in the cardinality of the input set. It is a self-referential set-to-sequence problem due to its input having no inherent order and the desired output being a permutation of the input elements. Given the same set of points as input, we expect the output to be the same tour, regardless of the order in which they are originally presented.

Further difficulties arise when we consider how to formalize the resulting reordering. One possible method involves the use of the aforementioned permutation matrices, which are discrete and therefore do not lend themselves to direct use of gradient-based backpropagation without a relaxation of the concept (Emami & Ranka, 2018). This and other formulations of representing a reordering (mentioned below) are discussed in more detail in section 4.



Figure 2: Travelling Salesman Problem

An example of a TSP input is given as a set of points in a 2D Euclidean space (top left) and an arbitrarily ordered array (bottom left). Individual dots represent x, y coordinates. As output, we see the shortest path between points (top right) and an array representing their target order (bottom right).

Such difficulties have motivated researchers to instead investigate differentiable, attention-based methods that involve pointing to the elements of the original set to define their permuted order. This approach often involves potentially computationally expensive beam search during inference.

Finally, if we forego the requirement of handling varying-length inputs, we can look to traditional learn-to-rank approaches for inspiration. In such frameworks the reordering is formulated as the assignment of a relevance score to each element, followed by sorting the elements according to that score, in monotonic order. However, sorting is a piecewise linear function, which therefore may contain many kinks where it is not differentiable. As a result, differentiable proxies to the sorting operator have been developed, but they did not achieve the expected $O(n \log n)$ time complexity until a method consisting of a projection onto a *permutahedron* was proposed by Blondel et al. (2020).

Alternatively, in learn-to-rank, our model may be trained to return *ranks*, i.e. positions of the input elements in the target (properly ordered) sequence. These ranks are piecewise constant functions, with derivatives that are either null or undefined, preventing gradient-based learning. However, significant progress has been made towards directly approximating ranking metrics (Rolínek et al., 2020) and constructing differentiable sorting and ranking operators (Blondel et al., 2020). Additionally, Engilberge et al. (2019) propose a deep neural net which can act as a differentiable proxy for ranking, allowing the use of traditionally non-differentiable metrics such as Spearman’s rank-order correlation (Spearman, 1904) as loss functions.

2.3 ML on Sets and Combinatorial Optimization

Set-to-sequence combines techniques from the field of machine learning on sets and combinatorial optimization. The former covers research areas related to both set-input and set-output problems, of which set-to-sequence is only concerned with the first kind. The latter consists of finding an optimal object from a finite set of objects and is strongly related to many forms of ordering tasks. The canonical example is the aforementioned TSP, which in itself has a long history of attempts at solving it through the most popular machine learning methods of the time, for example Smith (1999), Pihera and Musliu (2014), Ishaya et al. (2019) and Bengio et al. (2020).

Combinatorial optimization as such is of vital importance to modern industry applications. Consider the archetypal *Vehicle Routing Problem* (VRP), which poses the task of finding an optimal set of routes for a fleet of vehicles aiming to deliver goods to a given set of locations. The quality of the solution is determined by the *global transportation cost*. In the simplest variant of VRP, this is dependent on the sum of the lengths of tours for all vehicles. This effectively requires an ordering

of the locations into optimal trips, per each vehicle. Given the scale of modern logistical challenges and the environmental impact of freight, it is understandable that there have been many attempts to apply recent machine learning developments to such problems (Ibrahim, R, & Ishaya, 2019).

Current state-of-the-art combinatorial optimization algorithms often rely on handcrafted and hard-to-maintain heuristics for making decisions that are otherwise computationally infeasible or not well defined mathematically, for example Bello et al. (2016). It is a natural area of application for machine learning research and has been approached through the use of graph-based methods (Dai et al., 2017), reinforcement learning (Nazari et al., 2018) and attention mechanisms (Kool et al., 2019). For a comprehensive survey of the wider intersection of combinatorial optimization and machine learning, see Bengio et al. (2020).

2.4 Notation

The paper follows the notational conventions that are most common in literature. Scalar values are marked with lower case italics x_i , vectors with lower case bold typeface \mathbf{x} , matrices with capital case italics X . These matrices may be used to represent sets, in which case they are presented through roster notation with curly brackets, for example $X = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$.

However, sets may also consist of scalar elements, in which case a capital letter is still used to represent them: $A = \{a_i, \dots, a_n\}$. Given the importance of differentiating between unordered sets and ordered sequences, the latter are represented through angled brackets $\mathbf{x} = \langle x_1, \dots, x_n \rangle$ for additional clarity. When indicating the index of an element within a vector, whose symbol already contains a subscript (e.g. \mathbf{v}_j) the index of the scalar element is given in the superscript (v_j^i).

Individual permutations are marked as π , such that an example $\pi_i = \langle 3, 2, 1 \rangle$, consisting of integer indices referring to the original sequence $\mathbf{x} = \langle x_1, x_2, x_3 \rangle$, would result in the reordered sequence $\mathbf{x}^\pi = \langle x_3, x_2, x_1 \rangle$. In some cases, if the order of elements in the original sequence is nontrivial, a permutation π can also be given in two-line notation making both \mathbf{x}^π and the integer indices explicit:

$$\pi = \begin{pmatrix} x_3 & x_2 & x_1 \\ 3 & 2 & 1 \end{pmatrix} \quad (1)$$

2.5 Connections with Other ML Fields

In this section, a brief overview of other related topics from different fields of machine learning research is given. The aim is to point the reader who may only be tangentially concerned with set-to-sequence tasks to the appropriate area within their main field of interest. A reader with a decided focus on set-to-sequence is encouraged to continue reading section 3 directly.

2.5.1 NATURAL LANGUAGE PROCESSING (NLP)

There is a number of cases from the field of *Natural Language Processing* (NLP) that require tackling similar challenges to the ones faced by set-to-sequence methods. The popular sequence-to-sequence, encoder-decoder framework proposed by Sutskever et al. (2014) can itself, in principle, be applied to set-to-sequence problems, but does not perform well in practice. An example of such a case is the work on word ordering tasks, also known as *linearizations*, towards syntactically plausible word representations (Nishida & Nakayama, 2017).

The authors use sentences of ordered words to train the network to output a binary permutation matrix. When the original input, in the form of randomly ordered words from the target sentence, gets matrix-multiplied by this permutation matrix, the proper order is recovered. The network, referred to by the authors as the *Word Ordering Network* (WON), is an example of one way to formalize an ordering task. It can be seen as a simplification of the *Pointer Network* encoding method (Vinyals et al., 2015), discussed in more depth in a later section, whilst more closely resembling

the classic sequence-to-sequence models in the decoder (where it sequentially outputs rows of the permutation matrix).

Another example of an NLP task where permuting plays a key role is sentence ordering and order discrimination. The goal is to take a set of sentences and order them back into the original paragraph. Historically, this area of research has been dominated by hierarchical RNN-based approaches, which make use of LSTMs or GRUs in an auto-encoder framework (Logeswaran et al., 2018). First, a recurrent network is used to obtain the embedding of each sentence and then another to obtain the context representation of the entire paragraph.

More recent developments have seen the use of attention mechanisms to make it easier for this embedding to encode vital information regardless of the distance between information-carrying elements, as per the vanishing gradient problem. An example of this can be found in the *ATTOrderNet* architecture (Cui et al., 2018) and more recently in the *Set Transformer* (Lee et al., 2019), discussed in detail in section 3.5.

An example of an NLP set-to-sequence challenge where the output is not a permutation is the task of summarizing multiple documents into a single sequence of text as seen in the works of Ma et al. (2016) and Mani et al. (2018).

2.5.2 RANKING, INFORMATION RETRIEVAL AND CONTENT ORDERING

Another related, but succinctly distinct field comes in the form of ranking problems, information retrieval and content ordering. These encompass a family of challenges where there exists an *optimal* hierarchical order to the input elements, such that given two elements there is always a proper way in which they should be placed in relation to one another, which does not change depending on the other input-set elements.

More specifically, in the context of ranking for search, if we enter the query ‘cats’, the returned image of a cat should always rank higher than an image of a ‘dog’, regardless of what the other returned images may contain. This assumption is not always true in complex set-to-sequence problems, where any new element of the input-set can change what the proper relative order or structure of the already available elements’ sequence should be.

Traditionally, learn-to-rank problems have been tackled in a pairwise manner (Cohen et al., 1998), later approaches have applied neural methods on a list-based formulation of this problem (Cao et al., 2007). Ranking has also found application in content selection (Puduppully et al., 2019) and been employed as a useful auxiliary objective in a multitask setting for regression problems (Liu et al., 2019). A detailed look at listwise ranking approaches to ordering and structure prediction can be found in section 4.4.

2.5.3 SET REGRESSION, CLASSIFICATION AND SEGMENTATION

A more closely related area of work stems from set-input problems that have an output that is not a sequence. These include set regression, classification and segmentation challenges, among others. Effectively, this research field requires solving a near identical challenge to the first of two primary set-to-sequence challenges outlined at the beginning of this section, in that obtaining a proper encoding of the input set is vital.

Examples of such methods include *PointNets* for 3-dimensional point cloud classification and segmentation (Qi et al., 2017), which builds on previous work by Vinyals et al. (2016) in a specific geometric setting requiring both rotation and translation invariance (see section 3.6.3). Another example comes in the form of techniques for labelling objects based on a set of images from multiple viewpoints such as security cameras (Zhao et al., 2019a) and even fully convolutional models for set segmentation (Oliveira et al., 2020).

Methods that obtain the input set representation in a way that is interesting to set-to-sequence problems are included in the main section of this review and given appropriate focus, regardless of whether their original application was in sequence-output challenges.

2.5.4 SET-OUTPUT TASKS, INCLUDING SET-TO-SET

Conversely, a large field of work revolves around methods that learn to generate or predict a set as their output. Sets are the natural representation for many kinds of output data in machine learning tasks. These include a set of objects present in an image in an object detection setting (Zhao et al., 2019b), a group of points in a point cloud (Achlioptas et al., 2018) and a selection of nodes in a molecular graph for the problem of molecule generation (De Cao & Kipf, 2018).

The main challenge of set-output methods mirrors the primary challenge of representing sets in a permutation invariant way in set-input problems. If the order in which the model outputs elements does not matter, there are $n!$ equivalent, correct outputs that the model has to learn to consider equally good.

For example, imagine a simple task where the model must learn to take as input a set of integers and return the set of all primes present in the input. Specifically, given the input $A = \{1, 2, 3, 4, 5\}$ the correct output should take the form of the set $B = \{2, 3, 5\}$. However, given that ML model implementations operate on ordered, multidimensional arrays, the model must learn to treat all of these possible output sequences as equally correct: $\langle 2, 3, 5 \rangle$, $\langle 2, 5, 3 \rangle$, $\langle 3, 2, 5 \rangle$, $\langle 3, 5, 2 \rangle$, $\langle 5, 2, 3 \rangle$, $\langle 5, 3, 2 \rangle$.

Failure to properly account for this property of output sets leads to discontinuities that are difficult for most modern neural architectures to learn, even on seemingly trivial synthetic datasets (Zhang et al., 2020). An example that illustrates this comes in the form of an autoencoder trained to embed and then reconstruct input consisting of a set of n 2-dimensional points forming a regular polygon.

Every example in this dataset is a rotation of the same polygon around the origin. The discontinuity arises from this rotation, which forces a switch with regards to which element of the input set the model’s output neurons will be responsible for decoding. This is referred to as the *responsibility problem*. Proper handling of the set structure in the output requires the application of permutation invariant and permutation equivariant operations, much like in set-input problems, where the responsibility problem is not present.

Notable recent methods in the field of set prediction include the *Deep Set Prediction Network* (DSPN) by Zhang et al. (2019), which consists of a deep learning vector-to-set model that enables permutation invariance, and the *Transformer Set Prediction Network* (TSPN) by Kosiorek et al. (2020), that additionally takes advantage of the multiheaded self-attention introduced by Vaswani et al. (2017). The TSPN addresses the limitations of the DSPN related to set-cardinality learning. Additionally, an iterative attention mechanism referred to as *Slot Attention* has been proposed by Locatello et al. (2020), which decomposes input features into a set of representations, lending itself to set prediction tasks.

A sub-field of interest within set prediction is referred to as set-to-set, where both the input and output are structured as a set. Such tasks include recommendation (Sarwar et al., 2001), image search (Wang et al., 2014) and person re-identification (Zheng et al., 2015). Set-to-set challenges require both permutation invariance in the parts of the model that encode the input set and a proper cross-similarity function for the output sets (Saito et al., 2019), circumnavigating the responsibility problem.

2.5.5 ENSURING OTHER TYPES OF INVARIANCE

Finally, it may be of value to mention methods that obtain types of invariance other than under permutation. These methods stem from areas of application where the model’s final prediction should not be dependent on such predefined transformations. Examples include translational and rotational (also known as viewpoint) invariance, common in computer vision problems, addressed in the works of Ling et al. (2016) and Marcos (2016). Taking as illustrative case the task of object detection, to ensure the former quality the model needs to recognize the same object regardless of its position within the input image. To ensure the latter, given three dimensional images of a scene, the model must correctly identify an object regardless of the angle from which it is being perceived.

An important traditional approach to learning models that are invariant to certain transformation is data augmentation, as seen in Taylor and Nitschke (2018) and Hernández-García et al. (2019). Here, the underlying idea is that the model will be able to learn to become invariant to the chosen transformations once we augment our data with instructive examples that have been appropriately translated, rotated, partially obscured, blurred, illuminated, resized or had other applicable transformation applied to them, often in combination. More recent approaches that aim to safeguard machine learning models against relying on unintended aspects of data through adversarial strategies have also been proposed (Jaiswal et al., 2018). In order to apply the data augmentation approach to the set-to-sequence domain, we would increase our training set up to $n!$ times, providing the model with every possible permutation of each example set.

To prevent the costs associated with a larger training set, machine learning methods commonly employ various pooling operators after the stacked, equivariant feature extraction layers to obtain the desired invariance. However, in the case of translation invariance through pooled convolutional operations, the assumption that this completely prevents the model from exploiting the absolute location of an object in an image has been challenged (Kayhan & Gemert, 2020).

3. Set Encoding Methods

In this section, a qualitative comparison of different set encoding ML methods is provided, followed by detailed subsections devoted to the individual model architectures.

3.1 Method Comparison

This section introduces the reader to each of the relevant set-encoding methods in turn. These are also sometimes referred to in literature as *set-pooling* methods (Lee et al., 2019). Additionally, a number of comparison tables provides a summary overview: 3.1, 3.1.1.

It is important to note that some of the methods discussed in this section were designed specifically to handle set-to-sequence problems. As such, they contain both a set encoding module and a permutation outputting mechanism. Since it is not always immediately obvious how to combine a method that handles set encoding with a method that is designed to output a reordering, we compare various aspects of these methods in multiple places.

The models are compared with regards to the following aspects:

1. **Permutation Invariance:** whether the model obtains a permutation invariant representation of the input set. The same set must result in the same embedded representation, regardless of how the actual input array was permuted.

This feature does not guarantee that the final output of the model will be the same for differently ordered sequences obtained from a single set, as that may depend on the ordering method applied to the obtained permutation invariant set embedding, in order to output a sequence. This stems from the fact that these methods may require the reintroduction of the information relating to the order of the original input array and refer to it directly when outputting a permutation - particularly pointer-based attention. For more information on this, see section 4.

2. **Multiset Input:** whether the model can distinguish between a given input set and certain examples of its corresponding multisets, with repeated elements.

For example, the average() pooling operator will not be able to distinguish between a set $X = \{1, 2, 3\}$ and a multiset $X' = \{1, 1, 2, 2, 3, 3\}$. The max() operator will similarly fail in the case of $X'' = \{1, 1, 2, 3\}$.

3. **Complexity:** how the model's computational complexity relates to the cardinality n of the input set, and possibly other hyperparameters specific to a given architecture.

4. **Applications:** selected prominent domains which the model has been successfully applied to. Further examples can be found in the later sections devoted to each model.

Additionally, the most prominent architectures can be split into RNN-based methods, namely the Pointer Network and Read-Process-and-Write model, and the more recent fully connected ones, primarily the foundational DeepSets method and the Set Transformer.

Alternatively, an interesting way to distinguish between them would be to consider methods that depend on a variation of the attention mechanism introduced by Bahdanau et al. (2015), such as the Pointer Network and the Set Transformer, and others that do not utilize it.

Model	Perm. Invariance	Multiset	Complexity
Pointer Network (2015)	No	No ¹	$O(n^2)$
Read-Process-and-Write (2016)	Yes	Yes	$O(n^2)$
PointNet (2017)	Yes	No ²	$O(n)$
DeepSets (2017)	Yes	Yes ³	$O(n)$
Janossy Pooling (2019)	Yes	Yes	$O(n!)$
Set Transformer (2019)	Yes	Yes	$O(n^2)$
AttSets (2020)	Yes	Yes	$O(n)$
FSPooling (2020)	Yes	Yes	$O(n \log^2 n)$
RepSet (2020)	Yes	Yes	$O(mn + n^2 \log n)$

Table 1: Comparison of set encoding methods, part 1

3.1.1 NOTES ON COMPLEXITY

The complexity of the Read-Process-and-Write method is additionally impacted by the number t of steps in the *Process* block that computes the permutation invariant embedding of the input set. t is constant, but an interesting area of further research would be to learn it adaptively, similar to the method described by Graves (2016). The complexity of PointNet becomes $O(n^2)$ for 2D images and $O(n^3)$ for voxels, due to the convolutional operations. The authors of Janossy Pooling propose 3 methods of balancing tractability and the model’s representational power, as outlined in section: 3.6.2

Regarding the Set Transformer, the use of l stacked SAB layers results in quadratic complexity of $O(n^2l)$, use of a stack of l proposed ISAB layers with m inducing points results in complexity of $O(nlm)$. Similarly to the PointNet architecture, the complexity of AttSets grows depending on the dimensionality of the input, due to the use of convolutional layers in the encoder. However, the authors also provide a novel training paradigm, called *FASet*, and benchmark its mean time consumption for a single object against a selection of simple pooling methods, with favourable performance.

Regarding the RepSet method, m is the chosen number of hidden sets for the bipartite matching algorithm, represented by the columns of a trainable matrix. For more information on this and a proposed, more tractable relaxation, see section 3.6.5.

1. The Pointer Network does not treat its input properly as a set, therefore it cannot be said to properly handle multiset input either, but it will distinguish between input vectors with repeated elements.
2. Due to experimental results on the selected tasks, the authors of PointNet settle on $\max()$ as their recommended pooling operator, which does not distinguish between certain multiset variants. However, they report robust measurements of the performance of other pooling methods which can easily be included in the final model architecture and provide comparable results.
3. Depends on the pooling operator used after the stacked fully-connected layers, of which the authors of DeepSets primarily focus on $\text{sum}()$, which does distinguish between sets and multisets. However, $\max()$, which does not

Model	Applications
Pointer Network (2015)	combinatorial, multi-agent
Read-Process-and-Write (2016)	combinatorial, sorting
PointNet (2017)	3D shape classification and segmentation
DeepSets (2017)	set expansion, anomaly detection
Janossy Pooling (2019)	arithmetic, graph classification
Set Transformer (2019)	amortized clustering, anomaly detection
AttSets (2020)	3D shape reconstruction
FSPooling (2020)	set and graph classification
RepSet (2020)	text and graph classification

Table 2: Comparison of set encoding methods, part 2

3.1.2 NOTES ON DATASET PERFORMANCE

The listed set encoding methods can be applied to a wide spectrum of tasks. As a result, their performance has been tested on a variety of datasets, often in subtly different settings, which prevents direct comparison. In lieu of a table presenting their performance on a selected subset of such datasets, we provide a short discussion of the experimental results that do lend themselves to being compared. A more comprehensive experimental analysis in this area is a possible direction for future work.

Both Pointer Networks and the Read-Process-and-Write (RPW) method have been tested on the simple task of sorting a set of five floating point numbers between 0 and 1 (Vinyals et al., 2016). The Pointer Network achieved an accuracy of 0.90 compared to 0.94 reached by the RPW. Additionally, the RPW method appeared to be better at handling larger sets of floats. Both DeepSets and Janossy Pooling (Murphy et al., 2019) have been tested on simple arithmetic tasks such as sum-of-digits prediction, with each method reaching an accuracy of 1.0, albeit tested on input sets of different cardinalities. The Set Transformer has instead been tested on maximum value regression.

The Set Transformer, DeepSets and Janossy Pooling have also all been tested in terms of performance on unique count tasks. However, in the case of the Set Transformer experiments were performed on sets of handwritten characters from the Omniglot dataset (Lake et al., 2019), in the case of DeepSets on the MNIST8m hand-written digits (Loosli et al., 2007) and in the case of Janossy Pooling on simple integer sets.

The most commonly shared experimental task in the papers introducing and consequently utilizing the listed methods was point cloud classification. Particularly the ModelNet40 dataset (Wu et al., 2015) has been used to test four of the mentioned models. Whilst AttSets (Yang et al., 2020) employs it to formulate a multi-view reconstruction task, the other three methods are tested on the core classification task with PointNet reaching an accuracy of 0.892 (Qi et al., 2017), DeepSets 0.900 (Zaheer et al., 2017) and the Set Transformer 0.904 (Lee et al., 2019). However, the specific methods used to produce the point clouds from the provided mesh representation of objects showcased certain differences, further highlighting the need for a systematic, uniform comparison.

Both the Set Transformer and DeepSets methods have been tested on the task of set anomaly detection, specifically by way of the CelebA dataset (Liu et al., 2015). However, the DeepSets model was tested in terms of accuracy (0.75) and the Set Transformer in terms of the area under receiver operating characteristic curve and area under precision-recall curve, preventing direct comparison.

is also proposed as a problem-dependent variation. The formal proof extending DeepSets to multiset inputs was given by Xu et al. (2019).

The FSPool technique’s performance has been compared to the Janossy Pooling method through a visual question answering task, employing the CLEVR (Johnson et al., 2017) dataset. The accuracy of the latter was reported as 0.97 ± 0.54 , and of the former as 0.9927 ± 0.18 (Zhang et al., 2020).

Another useful task for the purposes of performance comparison is document classification, where given a document, the input to the model is the set of embeddings of its terms. DeepSets, Set Transformer and RepSet have been directly compared through their performance in this regard on 8 separate datasets (Skianis et al., 2020), with the Set Transformer consistently outperforming DeepSets, and RepSet outperforming both of the aforementioned methods.

Finally, the performance of DeepSets, Set Transformer and RepSet has been compared on the task of graph classification through the 5 datasets proposed by Kersting et al. (2016). The classification accuracy of DeepSets on the MUTAG dataset was 0.862, Set Transformer’s was 0.877 and RepSet’s 0.886. However, on the arguably more difficult IMDB MULTI dataset the Set Transformer outperformed RepSet, reaching an accuracy of 0.502, compared to 0.499. For a full overview, see the paper by Skianis et al. (2020).

Further details regarding the performance and limitations of presented methods are available in the sections devoted to them individually (below).

3.2 Pointer Networks

The Pointer Network (Vinyals et al., 2015) is an encoder-decoder neural network architecture including a modified attention mechanism, which allows it to learn a target reordering of input elements. It is the first deep learning method capable of taking sets as input and learning a desired permutation, resulting in complex output sequences.

Pointer Networks were originally designed to tackle combinatorial optimization problems with varying input sizes, which was their main advantage over previous sequence-to-sequence methods. A Pointer Network can be trained on inputs of varying length and has been demonstrated to generalize reasonably well to unseen lengths (Vinyals et al., 2015).

Additionally, Pointer Networks included a modification of the content-based attention mechanism introduced by Bahdanau et al. (2015) which made it possible to treat the output of the model as pointers to elements of the input sequence. This attention-based pointing is one of the most popular methods for giving models the ability to output a permutation of the original input, regardless of the way they encode the original set. Due to its importance as a purely element-ordering technique, it is separately described in further detail in section 4.2.

3.2.1 POINTER NETWORKS LIMITATIONS

An important characteristic of Pointer Networks is that they do not strictly treat the input as a set, instead processing it solely through sequential recurrent neural networks. As a partial consequence they do not obtain a permutation invariant representation of the encoded set. This results in a situation where the same set can be represented as two differently ordered input arrays, leading to the model predicting two different outputs for it. Thus returning the optimal order is not guaranteed.

Another important limitation is that nothing is explicitly preventing the model from outputting an invalid reordering of the input set or sequence. This becomes apparent during early training, when the model points to the same elements of the input at various indices of the output sequence. However, this can be mitigated by the addition of a beam search mechanism to the decoder during inference or by progressive masking. In the latter case, the entry in the attention vector referring to an element that had already been pointed to is preset to an infinitely negative value at each successive iteration, preventing it from being pointed to again, at the cost of certain inductive bias being introduced into the model.

3.2.2 POINTER NETWORKS DETAILS

The Pointer Network consists of a recurrent neural network (RNN) encoder and an RNN decoder with a modified attention mechanism. The model obtains a content-based attention vector $\mathbf{a}_j \in \mathbb{R}^n$ at each decoder step j . This vector represents the conditional probability of each input element x_i being the correct one to be pointed to at this step, conditioned on all previous steps as well as the entire input sequence $\mathbf{x} = \langle x_1, \dots, x_n \rangle$, in the form of all encoder hidden states $E = \langle \mathbf{e}_1, \dots, \mathbf{e}_n \rangle$ obtained when the encoder block iterates over the input array.

For simplicity, we will assume that each element x_i must be pointed to exactly once, meaning that the output sequence of nonnegative integer pointers $\mathbf{y} = \langle y_1, \dots, y_n \rangle \in \mathbb{Z}^n$ represents a valid permutation π , such that a sample target output $\mathbf{y}^\pi = \langle 0, 2, 1 \rangle$ would represent the reordered sequence $\mathbf{x}^\pi = \langle x_1, x_2, x_3 \rangle$ for the sample input $\mathbf{x} = \langle x_1, x_3, x_2 \rangle$. This will mean that when iterating over both encoder states \mathbf{e}_i and decoder states \mathbf{d}_j , we will always be in range 1 to n . The input sequence \mathbf{x} can itself consist of multidimensional elements, or such embeddings of each x_i can be obtained prior to the pointer network module through a chosen embedding layer.

The attention mechanism in the decoder block is as follows:

$$z_j^i = \mathbf{v}^T \tanh(W_1 \mathbf{e}_i + W_2 \mathbf{d}_j) \quad \text{for } i \in (1, \dots, n) \quad (2)$$

$$\mathbf{a}_j = \text{softmax}(\mathbf{z}_j) \quad \text{for } j \in (1, \dots, n) \quad (3)$$

$$P(y_i | y_1, \dots, y_{i-1}, \mathbf{x}) = \mathbf{a}_j \quad \text{for } j \in (1, \dots, n) \quad (4)$$

Where \mathbf{d}_j is the decoder's hidden state at the j -th output element, \mathbf{e}_i is the encoder hidden state at the i -th input element, W_1 , W_2 and \mathbf{v} are trainable tensors. The \mathbf{z}_j vector is of the same length as the input \mathbf{x} and represents an output distribution over the dictionary of input elements. After the application of the softmax nonlinear activation function, turning it into \mathbf{a}_j , it becomes an attention vector. For a visual explanation, see Figure 3.

3.2.3 POINTER NETWORKS APPLICATIONS

The primary application of Pointer Networks are tasks where the target output is a reordering of the elements of the initial input. This reordering is based on pointers to indices of the original input sequence. Examples of such problems in currently active research areas include element sorting, coherence modeling (Logeswaran et al., 2018), word ordering (Cui et al., 2018) and sentence ordering (Wang & Wan, 2019), as well as summarization (Sun et al., 2019) and ranking in information extraction (Bello et al., 2018).

In the original paper, the Pointer Network models have been tested on challenging combinatorial optimization problems such as finding planar convex hulls, computing Delaunay triangulations and the Travelling Salesman Problem. Experiments have shown that even with computationally intractable, NP-Hard problems such as TSP, this model architecture was able to learn competitive approximate solutions, limited by the scale of the problem, with $n \leq 50$ for the TSP.

Pointer Networks also found usage within the AlphaStar reinforcement learning model which defeated a grandmaster level player in the competitive real-time strategy game of Starcraft II (Vinyals et al., 2019). They were employed to help the agent manage the structured, combinatorial action space in conjunction with an auto-regressive policy.

3.3 Read-Process-and-Write Model

The Read-Process-and-Write (RPW) model is a neural network architecture consisting of three distinct blocks and aiming to obtain a permutation invariant representation of the input set, whilst learning a function mapping it to arbitrary target outputs (Vinyals et al., 2016).

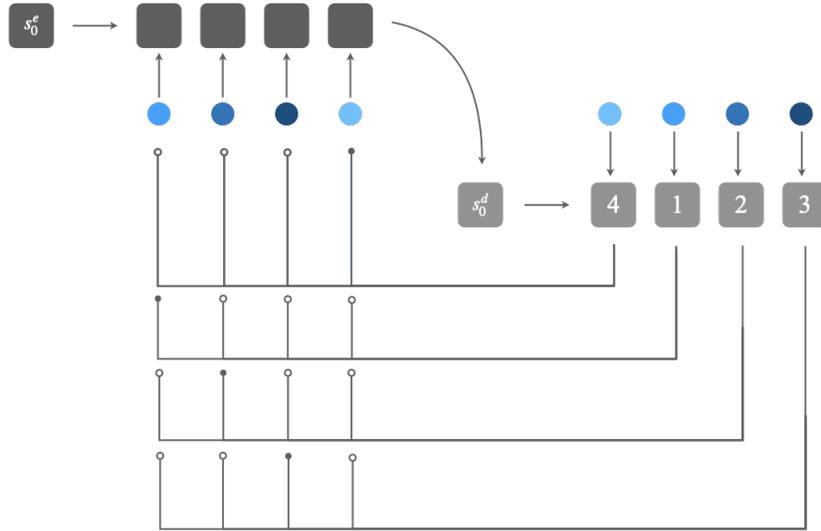


Figure 3: Pointer Network

In Figure 3, an encoding RNN sequentially processes each element of the input array (blue dots), encoding it into a hidden state s_n^e (dark gray), which is fed to the pointing decoder RNN (light gray). At every step, the second network produces a vector that represents a content-based pointer attention over the encoded inputs.

RPW satisfies the key property of obtaining a permutation invariant representation of its input through a variation of the attention mechanism. It can be seen as a special case of a *Neural Turing Machine* (Graves et al., 2014) or a *Memory Network* (Weston et al., 2015) in that it is a recurrent neural network model that creates a memory representation of each element in the input sequence and accesses it via an attention mechanism.

In the original RPW paper its authors also demonstrated that the order in which elements are organized as input has a significant impact on the learning performance of earlier sequence-to-sequence architectures. This is an important observation given the fact that the recurrent neural networks employed in them are, in theory, universal approximators (Schafer, 2007).

3.3.1 RPW LIMITATIONS

Whilst the RPW model constitutes a significant improvement in the way machine learning methods handle input sets, it suffers from the same limitation as Pointer Networks in terms of ordering their elements into the output sequence. Namely, it is not strictly prevented from pointing to the same element of the input set multiple times in the output, effectively returning either an invalid sequence or an incomplete reordering. This is a particularly important limitation in relation to handling *multisets* (also known as *msets* or *bags*), where the same element can occur multiple times in the input. However, it can be mitigated through beam search or progressive masking as described in the Pointer Network section. It also suffers from a significant decrease in performance as the size of the input set increases.

3.3.2 RPW DETAILS

The RPW architecture consists of three distinct blocks:

1. *Read Block* - which embeds every element of the input set using the same neural network for each $x_i \in X$.
2. *Process Block* - which consists of a recurrent neural network that evolves its hidden state using a modified content-based attention mechanism to obtain a permutation invariant representation of the input over a separately predefined number of steps t .
3. *Write Block* - which takes the form of a Pointer Network in the set-to-sequence tasks but can also be another recurrent neural network decoder for tasks where the output elements come from a fixed dictionary.

The Process Block evolves the permutation invariant representation of the input set by repeating the following steps t times:

$$\mathbf{q}_t = \text{LSTM}(\mathbf{q}_{t-1}^*) \quad (5)$$

$$z_t^i = f(\mathbf{m}_i, \mathbf{q}_t) \quad (6)$$

$$a_t^i = \frac{\exp(z_t^i)}{\sum_j \exp(z_t^j)} \quad (7)$$

$$\mathbf{r}_t = \sum_i a_t^i \mathbf{m}_i \quad (8)$$

$$\mathbf{q}_t^* = \langle \mathbf{q}_t, \mathbf{r}_t \rangle \quad (9)$$

where i is the index over all embedded elements of the memory vector \mathbf{m}_i obtained by the Read Block, \mathbf{q}_t is effectively a query vector allowing us to read the permutation invariant representation \mathbf{r}_t from the memories using an attention mechanism and $f()$ is any differentiable operation that takes two vectors and returns a scalar, most commonly the dot product $f(\mathbf{a}, \mathbf{b}) = \mathbf{a} \cdot \mathbf{b} = \sum_{i=1}^n a_i b_i$.

An important implementation nuance is related to the third step, where the attention vector \mathbf{a}_t is obtained via a softmax operation. Depending on weights initialization, that step can result in the undefined operation of dividing infinity by infinity. This can be prevented by bounding the value range of the \mathbf{z}_t vector by the use of the tanh function (Logeswaran et al., 2018).

The LSTM() is a recurrent neural network that takes no inputs, only evolving the hidden state \mathbf{q}_t . The final set encoding \mathbf{q}_t^* , is obtained by concatenating the previous hidden state \mathbf{q}_t and the permutation invariant representation \mathbf{r}_t . \mathbf{q}_t^* becomes the hidden state input during the next iteration t of the Process Block.

3.3.3 RPW APPLICATIONS

The RPW architecture has been applied to both continuous and discrete inputs. In the former case, the input can be a set of floating point numbers or a high-dimensional embedding of the entities of interest. In the latter, it can be dictionary entries. Considerations related to the specific structure of those input elements are out of scope for this paper as they pertain to the wider topic of representation learning.

This model architecture has also been used in few-shot object recognition (Xu et al., 2017), graph classification (Ying et al., 2018) and one-shot learning in the context of drug discovery (Altae-Tran et al., 2017). The original paper tests it on the problem of sorting a varying-size set of floating point numbers between 0 and 1. It achieves accuracy of 94% on sets of 5 elements, performance dropping significantly for larger ones (50% for 10 elements, 10% for 15).

This key property of obtaining a permutation invariant representation was further formalized in the DeepSets paper, presented in the following section.

3.4 DeepSets

The DeepSets framework (Zaheer et al., 2017) provides a robust mathematical analysis for designing *permutation invariant* and *permutation equivariant* deep learning models. Both of these concepts are explained in section 2.1. Where the Vinyals et al. (2016) paper focused on the former in the setting of a specific practical application, the authors of DeepSets provide a generic framework for the proper handling of set-inputs for both supervised and unsupervised learning.

The primary contribution of the DeepSets method lies in tackling the first part of the more general set-to-sequence challenge, which is to implement arbitrary set functions that result in permutation invariant representations. We have already described one example of an operation that ensured this in the RPW model’s Process Block, namely a variation of the content-based attention mechanism consisting of a modified recurrent neural network. DeepSets propose a simpler, sum-based method to achieve this.

The DeepSets framework offers a simplified procedure by relying on summation of all element representations prior to further nonlinear transformations, which then transform these summed representations into the desired output (e.g. a class probability distribution or a single number for set regression). Later reimplementations of the DeepSets architecture experiment with replacing the sum operation with other permutation invariant alternatives such as taking the mean or maximum, with comparable results (Lee et al., 2019).

Additionally, the DeepSets analysis expands upon the set-input challenge by allowing for permutation equivariance, where the order of the output elements mirrors the order of the input sequence. This can be achieved, in one case, by adding a diagonal symmetry and diagonal identity constraint to the weights matrix of a fully-connected neural network layer, prior to the nonlinearity. However, set-to-sequence methods do not make extensive use of permutation equivariance therefore these are not detailed here. For more information, see the original paper.

3.4.1 DEEP SETS DETAILS

The proposed permutation invariant function for inference over sets takes the following general form:

1. Each element x_i in the input set is transformed *independently* into an embedded representation $\phi(x_i)$, possibly through multiple layers of a feed-forward neural network.
2. The representations $\phi(x_i)$ are summed together and the result is further processed using another network ρ consisting of any number of fully-connected layers with nonlinearities.

Both ϕ and ρ can be replaced by universal approximators, which can be learned to approximate arbitrary polynomials. In cases where additional information q is available, it can be used to obtain the conditional mapping $\phi(x_i|q)$. The key to permutation invariance in this framework is simply summation of the obtained per-element representations.

$$\text{DeepSets}(\{x_1, \dots, x_n\}) = \rho(\text{sum}(\{\phi(x_1), \dots, \phi(x_n)\})) \quad (10)$$

3.4.2 DEEP SETS LIMITATIONS

This approach is much simpler to implement, compared to the RNN-based Pointer Networks and the RPW model. However, it generally prevents the model from learning pair-wise and higher order interactions between the elements of the set, which are lost during the summation. Additionally, significant doubts have been raised by Wagstaff et al. (2019) relating to the limits of the representational power of the DeepSets method. More precisely, the $O(n)$ computational complexity comes at the cost of the dimensionality of the latent space having to be at least equal to the cardinality of the input set n to ensure universal function approximation.

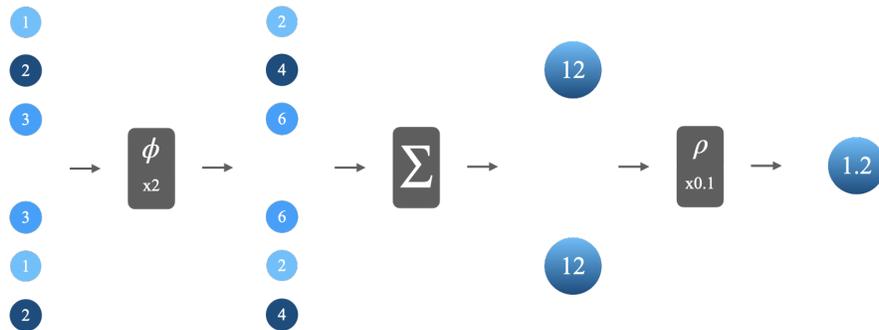


Figure 4: DeepSets

Every element of the two identical, shuffled sets of blue dots (leftmost) is embedded in an independent and identical way by the ϕ layer, resulting in a permutation equivariant transformation. These are then summed (Σ) into a permutation invariant representation and further transformed by the ρ layer.

3.4.3 DEEP SETS APPLICATIONS

The DeepSets framework has been applied to point cloud classification (Qi et al., 2017), generalization tasks in reinforcement learning (Karch et al., 2020), outlier detection and anomaly classification (Oladosu et al., 2020) among others.

3.5 Set Transformer

One of the most elaborate methods designed for set-input problems is the *Set Transformer* (Lee et al., 2019). This method can be seen as an extension of the popular feed-forward, attention-based Transformer (Vaswani et al., 2017) to the domain of machine learning on sets.

The Set Transformer consists of the expected stacked multi-head self-attention layers for both the internal encoder and decoder as seen in the classic Transformer. One aspect that separates it from the previously described set-to-sequence methods is that instead of using a fixed pooling operation such as summing or taking the average to ensure permutation invariance, it employs a parameterized pooling function that is learned and therefore much more adaptive to the particular task at hand. This is further referred to as *Pooling by Multihead Attention* (PMA) and explained in more detail later in this section.

The Set Transformer is specifically designed to model higher-order interactions among elements and their subsets within the input set, whilst satisfying the permutation invariance and variable input size requirements common to set-to-sequence problems. Its key novel contribution is that it concurrently encodes the entire input set through a sequence of permutation equivariant *Set Attention Blocks* (SABs). By comparison, the previously discussed DeepSets method obtained element features independently of other input set elements. This modification allows the Set Transformer to explicitly learn pairwise and even more complex interactions between set elements during the encoding step, dependent on the number of stacked SAB layers.

3.5.1 SET TRANSFORMER LIMITATIONS

However, the Set Transformer also introduces certain costs. The SAB, a proposed variation on the multihead attention block employed in the classic Transformer, which enables the Set Transformer to encode higher-order interactions between set elements, has the limiting quality of requiring quadratic time complexity $O(n^2)$ relative to the cardinality of the input set n .

The authors of the method address this limitation by proposing an *Induced Set Attention Block* (ISAB). It takes advantage of a vector of *inducing points*, which is of predetermined size and is used to obtain a hidden representation of the input set by attending to it. This is in effect a low-rank projection that might be familiar to readers who have experience with autoencoder models. The technique reduces the required computation time to a linear $O(mn)$, where m is the chosen number of inducing points and n is the set cardinality, at the cost of reduced performance and an additional hyperparameter to be tuned.

3.5.2 SET TRANSFORMER DETAILS

What follows is a more detailed inspection of the SAB, ISAB and PMA methods, which assumes some familiarity with the multihead attention mechanism proposed by (Vaswani et al., 2017) as per the Transformer architecture. However, to establish full clarity, we will formally define both the basic Transformer attention and its multihead variant first.

Assume we have a set of n elements as our input, each with dimensions d_e . First, we obtain three new vectors per each element by multiplying it with three matrices, whose weights are learned during the training step. These transformed input representations are further referred to as n query vectors $Q \in \mathbb{R}^{n \times d_q}$, key vectors ($K \in \mathbb{R}^{n \times d_k}$) and value vectors ($V \in \mathbb{R}^{n \times d_v}$). These are then mapped to the desired attention outputs, applying an activation function such as softmax in the following way:

$$\text{TransformerAttention}(Q, K, V) = \text{softmax}(QK^\top)V \tag{11}$$

The original implementation includes a scaling factor, which has been omitted from the above equation for the sake of simplicity. The pairwise dot product of the query Q and key K vectors measures how related each pair is. The final output is a weighted sum of V .

The multihead attention mechanism extends this further by projecting each query, key and value onto h separate vectors via sets of the three parameter matrices W_i^Q, W_i^K, W_i^V , one per each of the h heads. Then, the $\text{TransformerAttention}()$ function is applied to each of these h vectors to obtain each head’s preliminary output O_i :

$$O_i = \text{TransformerAttention}(QW_i^Q, KW_i^K, VW_i^V) \tag{12}$$

Finally, these outputs $\{O_i\}_{i=1}^h$ are concatenated and then linearly transformed:

$$\text{MultiheadAttention}(Q, K, V) = \text{concatenate}(O_1, \dots, O_h)W^0 \tag{13}$$

Now we can begin to move on to the Set Attention Block (SAB). It is designed to take a set and perform a slightly modified self-attention operation between its individual elements, which results in an output set of the same size. It will be useful to first define a *Multihead Attention Block* (MAB), which is an intermediate building block in both SAB and its less computationally intensive alternative - ISAB. MAB takes as input two matrices of the same dimensions: $A, B \in \mathbb{R}^{n \times d}$ and first obtains a hidden representation Z using a layer normalization operation as defined by Ba et al. (2016):

$$Z = \text{LayerNormalization}(A + \text{MultiheadAttention}(A, B, B)) \tag{14}$$

Then, it processes each element in each row of Z in an independent, identical way, in the same manner we have seen as part of the DeepSets method, prior to the summation. This can be done via a row-wise feed forward layer $\phi()$:

$$\text{MAB}(A, B) = \text{LayerNormalization}(Z + \phi(Z)) \tag{15}$$

However, in our set-to-sequence setting we do not have two separate matrices of sets A and B that we wish to encode into some joint representation. Therefore the actual SAB() is defined on a single matrix of an input set $X \in \mathbb{R}^{n \times d}$.

$$\text{SAB}(X) = \text{MAB}(X, X) \quad (16)$$

If we stack l SAB layers in our set encoder, which we are able to do since both the input and output of the Set Attention Block is a set of the same size, the model’s computational complexity in the cardinality of this set is $O(n^2l) \approx O(n^2)$. At the cost of quadratic computation time, stacking 2 SAB layers enables the model to encode pairwise interactions between elements. Stacking more such layers makes explicitly encoding higher-order interactions possible, which is a crucial novel contribution for tasks where such interactions define the target output.

The authors of the Set Transformer method address the computational cost of SAB by proposing a less expensive variation of it, called the Induced Set Attention Block. In ISAB, an additional array of *inducing points* $I \in \mathbb{R}^{m \times d}$ is included. This vector is of predefined dimension m , resulting in a computational complexity $O(mn)$ or $O(lmn)$, if stacked ISAB layers are applied. The calculations within ISAB are defined as follows:

$$\text{ISAB}_m(X) = \text{MAB}(X, \text{MAB}(I, X)) \quad (17)$$

The learned values of the inducing points I are expected to encode large-scale aspects of the input set as meaningful features for the ultimate task. Similar to SAB, the output of ISAB is a set of the same size as the input set, which is why we still need a pooling operation to be applied to it at this point.

The final aspect of the Set Transformer that distinguishes it from the earlier set-encoding methods is the Pooling by Multihead Attention (PMA) stage. Unlike a simple sum, mean or max, the PMA pooling function has learnable parameters, which allows it to increase or decrease the relative importance given to the encoding of individual encoded elements of the output of the SAB and ISAB blocks. Usage of PMA requires specifying the number k of seed vectors $S \in \mathbb{R}^{k \times d}$. Assuming we have already obtained the encoded set features $E \in \mathbb{R}^{n \times d}$ via stacked SAB or ISAB layers:

$$\text{PMA}_k(E) = \text{MAB}(S, \phi(E)) \quad (18)$$

In most cases a single ($k = 1$) seed vector is used, resulting in a single pooled set encoding, but certain clustering tasks may require multiple related outputs, justifying the use of a larger k .

3.5.3 SET TRANSFORMER APPLICATIONS

The overall Set Transformer architecture can, in principle, be applied to any set-input problem (and therefore any set-to-sequence task). It will perform particularly well in problems where pairwise and higher-order interactions between the input set’s elements are important to the task at hand. The authors of the original paper demonstrate the model’s usefulness within such areas on the challenge of counting unique characters in a set of input images, amortized clustering with mixture of Gaussians, point cloud classification and set anomaly detection.

It has since been applied to set-of-sets embedding problems (Meng et al., 2019) and transfer learning in dialogue systems (Wolf et al., 2019).

3.6 Other Set-Input Methods

Below is a list of other set encoding methods that do not necessarily lend themselves directly to set-to-sequence problems, but may be of interest to the reader. The specifics of using the learned, permutation invariant set representation to produce a sequence of set elements are discussed at the beginning of the next chapter, specifically in subsection 4.1.

3.6.1 FEATUREWISE SORT POOLING

This method, also known by its abbreviation as FSPool (Zhang et al., 2020), came from the field of set prediction, in relation to a problem where both the input and output can be conceived of as sets. The authors expand upon one of the naive approaches to encoding sets in a permutation invariant way. Namely, the technique of simply sorting all the elements of the input set by the values of their single chosen feature, as seen in previous work by Zhang et al. (2018).

In set-output problems this approach results in discontinuities in the optimisation whenever two elements swap positions after the sort. This is referred to in set prediction challenges as the *responsibility problem* (Zhang et al., 2020). To avoid this difficulty, the authors have developed a novel pooling method which sorts each feature across the elements of the input set and then performs a weighted sum.

This allows the model to remember the permutation applied through the featurewise sorting and apply its inverse in the decoder. This process restores the original, arbitrary order of the input elements making the encoding a permutation equivariant operation, preventing the discontinuity in the outputs of the model.

3.6.2 JANOSSY POOLING

Another interesting approach to set-encoding through the use of simpler pooling operators was proposed by Murphy et al. (2019). In the titular *Janossy Pooling*, the symmetric (permutation invariant) encoding function is expressed as the average of a mixture of permutation sensitive functions applied to all reorderings of the original input.

This approach immediately raises the question of tractability. Generating all permutations of a set results in $n!$ intermediate inputs, all of which would then require the application of the chosen permutation sensitive function. To mitigate this, the authors propose a number of strategies, among them the use of a smaller number of selected canonical orderings that are presumed to carry relevant information for the specific task at hand, such as simple sorting, *betweenness centrality* and others (Niepert et al., 2016).

As an alternative to canonical orderings, the authors also propose a method related to a model’s ability to explicitly learn pairwise and higher-order interactions between the elements of the input set. This method is referred to as k -ary dependencies. It consists of projecting the input to a length k sequence, for example by only keeping the first k elements, limiting the number of permutations that need to be averaged to $k!$, which can be tractable for a small enough k . The number k becomes a hyperparameter capable of balancing tractability with the model’s ability to learn k -ary interactions in the input. Finally, the authors also experiment with permutation sampling as the third method of reducing the computational complexity of Janossy Pooling, as proposed by Moore and Neville (2017) and Hamilton et al. (2017) in relation to machine learning on graphs.

3.6.3 POINTNET

Not to be confused with the Pointer Network described in section 3.2, *PointNet* is a set-encoding method designed to handle 3D point clouds proposed by Qi et al. (2017). The geometric setting of the problem tackled by the authors of this model shares many similarities with the set-to-sequence domain. These include the lack of order in the input, requiring the use of symmetric, permutation invariant encoding functions, and the importance of interactions between individual input elements or, specifically in this case, points. An additional requirement for this setting is invariance under certain geometric transformations of the entire point cloud. For example, if we are tasked with classifying 3D objects represented by a point cloud, we want to correctly classify a chair regardless of its rotation and translation.

In practice, PointNet first obtains an embedding of each of the input points through stacked, fully-connected layers in the form of a multilayer perceptron, such that each element is identically

and independently transformed. This permutation equivariant representation is then pooled via the $\max()$ operator (per dimension) and further transformed through an additional fully-connected layer. This is in effect a slight variation of the procedure proposed by Zaheer et al. (2017) and discussed in section 3.4. Other pooling operations, including an attention-based sum inspired by the RPW model (Vinyals et al., 2016), are also experimented with in the original paper..

Finally, the obtained point cloud encoding is concatenated with the embedding of each point, reminiscent of the approach seen in listwise ranking, described in section 4.4. This combination of local and global features is shown to be crucial for point segmentation tasks. The authors also provide proof that their network is a universal approximator for continuous set functions and demonstrate its robustness to small perturbations of the input set.

3.6.4 ATTSETS

The *AttSets* model, proposed by Yang et al. (2020), uses weighted attention to obtain a permutation invariant representation of the input set. It was originally applied to a multi-view 3D reconstruction task, where a set of images of the same object from different angles is used to estimate its true three dimensional shape.

AttSets improves the performance of previous, simpler pooling functions used for 3D object recognition. These include both first-order operators such as $\max()$, $\text{average}()$ and $\text{sum}()$, which do not have any trainable parameters, as well as higher-order statistical functions such as bilinear pooling (Lin et al., 2018), log-covariance pooling (Ionescu et al., 2015) and harmonized bilinear pooling (Yu et al., 2018), which have only few.

In order to achieve this, each element of the set is individually and independently transformed via a learned attention function, which can take the form of a fully connected layer or a multidimensional convolutional layer, depending on the form of the input. The output of this function is normalized via $\text{softmax}()$ and then used as an attention mask over the original input elements. This allows the model to learn to pay a varying degree of attention to individual dimensions of the input elements' representations. To obtain the final, fixed-length set encoding, the original input elements are multiplied by the attention mask and summed together.

3.6.5 REPSET

An interesting set-encoding method, referred to as *RepSet*, has been proposed by Skianis et al. (2020). The RepSet model consists of stacked feed-forward, fully connected layers, reminiscent of the DeepSets method (Zaheer et al., 2017), followed by a custom permutation invariant layer replacing the $\text{sum}()$ operator. This layer is inspired by concepts from the field of bipartite graph matching and has allowed the model to show promising performance on text and graph classification tasks.

The permutation invariance is achieved through a configurable number of hidden sets (potentially of different cardinalities), whose elements correspond to columns of trainable weight matrices. These are then compared with the elements of the actual input set to create matrices that are fed to a bipartite matching algorithm, specifically the Hungarian Algorithm (Grinman, 2015). The resulting values can be further transformed through standard neural network layers for set classification and regression purposes.

A significant issue with this approach is the computational complexity $O(mn + n^2 \log n)$, where n is the cardinality of the input set and m is the chosen number of hidden sets. This characteristic, limiting the usefulness of the method regarding larger sets, stems from the bipartite matching algorithm needed to obtain the final set encoding. The authors of the method address this by proposing a relaxation of RepSet, referred to as *ApproxRepSet* (Skianis et al., 2020), which removes one of the constraints on the range of values taken by the elements of the hidden sets.

4. Ordering Methods

This section focuses on the second of the two primary challenges inherent to set-to-sequence tasks, which is outputting a permutation. Assuming we are able to obtain a meaningful representation of an input set of any length, how do we use that representation to produce a reordering of the input's original elements? This is a constraint that is easy to satisfy when designing traditional combinatorial optimization algorithms, yet in deep learning it requires relatively complex model architectures (Bengio et al., 2020).

Three particular branches of ordering methods have emerged as most prominent in deep learning:

1. *Attention-based Pointing*

In which a vector of attention weights over all elements of the input set is generated iteratively at every index of the output sequence. The highest attention value points to the element that should be placed in the current position within the sequence.

2. *Permutation Matrices*

Where a square, binary, doubly-stochastic matrix or a relaxation thereof is generated for each input set. The index of the highest value in each row identifies the element that should take the position at the same index as the number of the row. The input can be left-multiplied by this matrix to obtain the final reordering.

3. *Listwise Ranking*

In ranking methods the target order is represented through the assignment of a score to each element of the input set, which enables the final permutation to be obtained through sorting. Listwise ranking takes into consideration the relative scores of all other set elements when computing the score for a particular one.

Each of these three basic frameworks is described in the following sections, with references to specific methods that expand upon them, where relevant. First, however, we must discuss how best to utilize the permutation invariant set representation, obtained via the set encoding methods discussed in section 3, to output the target sequence.

4.1 From Set Representation to a Sequence

As stated previously, deep learning models do not directly take unordered sets as input. Instead they transform ordered arrays representing one (usually arbitrary) of $n!$ permutations of a given set's elements. Therefore, an inductive bias is introduced into the model's internal architecture to first obtain a permutation invariant representation of the underlying input set, which by definition will be the same regardless of which arbitrary permutation the model happened to receive.

However, when our prediction target is the optimal order of a given set's elements and we are directly feeding our model an array in some arbitrary order, then our target output sequence must refer to this initial, random order when predicting the preferred one. In effect, in the case of simple permuted sequences without repetition and exclusion, we are seeking a permutation equivariant function that outputs ranks, conditioned on a permutation invariant representation of the entire set. In the case of complex permutations, we also expect the same sequence to be predicted based on the same input set.

Regardless, the output sequence has to refer to the arbitrary order in which the input set elements are presented to it. In essence, whilst we assume there is one optimal order of a given set's elements, there are $n!$ target permutations relative to the $n!$ arbitrary ones that form the actual input of the model, per input set. This is because each target permutation reorders an arbitrary one that the model receives as actual input, each representing the same, singular optimal order of the underlying set's elements.

The remaining question is how to utilize the fixed-length, permutation invariant set representation when outputting the sequence representing a permutation of the arbitrarily ordered input’s elements. There are two dominant approaches, namely summation and concatenation of the representation of the entire set with the learned representations of each element. In this way, every transformation of individual set elements has the chance to take into consideration the entire available set, as previously seen in the PointNet model by Qi et al. (2017).

Considerations regarding which of those two methods is preferable relate to the larger field of representation learning, with ample examples both on the side of summation (Szegedy et al., 2016) and concatenation (Noreen et al., 2020). These include the *ResNet* architecture by He et al. (2016), with addition used in the eponymous residual connections, and the *DenseNet* model and its descendants, employing concatenation (Huang et al., 2018a), within the field of computer vision.

An alternative approach, applicable to RNN-based ordering methods, would be to utilize the set representation in the hidden state of the recurrent network. Similarly, attention-based methods could use the set representation to influence which elements the model focuses on. Finally, in ranking approaches the set representation can rather elegantly be used in place of the *query* (see section 4.4 for more details). Investigation of the efficacy of these methods is a possible direction for future research, as outlined in section 5.6.

4.2 Attention-Based Pointing

The term *attention* covers a wide spectrum of methods within the field of machine learning. First introduced as part of an encoder-decoder model applied to a sequence-to-sequence translation task by Bahdanau et al. (2015), it has since been utilized in a number of other domains, including computer vision (Xu et al., 2015), graph-based problems (Veličković et al., 2018), reinforcement learning (Iqbal & Sha, 2019) and many more. For a comprehensive overview, see Chaudhari et al. (2019).

In structured output tasks that are inherent to the set-to-sequence domain, the key idea is to induce a vector of attention weights over the entire input, allowing for the selection of elements by their position, conditioned on the learned representation of the entire input and the sequentially predicted outputs. An example of such a setting from the field of natural language processing (NLP) is *linearization*, which involves solving a problem that will be familiar to most readers who have tried to learn a new language. Namely, given a randomly shuffled set of words, we are tasked with ordering them into a properly structured sentence. A survey of these and other attention-based methods in NLP can be found in the work of Hu (2019).

In effect, the model can learn to identify the current most pertinent position in the input through the attention weights, which we then interpret as pointers to the next preferred element for the output sequence. It may be worth mentioning that valid concerns have been raised as to the overall interpretability of attention weights in text classification and question answering tasks by Jain and Wallace (2019). However, in attention-based pointing these weights have a direct impact on the predicted permutation, leaving little room for interpretational ambiguity.

4.2.1 DETAILS OF ATTENTION-BASED POINTING

Attention-based pointing is an adaptation of content-based attention (Graves et al., 2014) to an ordering challenge. The model learns a distribution over all input elements at each position of the output sequence. This should not be confused with self-attention, also known as intra-attention (Cheng et al., 2016), where the target sequence is always the same as the input. In self-attention, the parameterized attention function learns to relate different positions of the input sequence to each other. In attention-based pointing the attention function learns to relate the elements of the input with the current positions of the output.

Assume we have a randomly ordered input sequence of varying length n , consisting of d -dimensional elements: $X = \langle \mathbf{x}_1, \dots, \mathbf{x}_n \rangle$ such that $X \in \mathbb{R}^{n \times d}$, which we want to order according to some pref-

erence. This sequence is an n -tuple, representing a single permutation of the corresponding set $X' = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$, which we cannot represent directly as a single vector, due to the fact that vectors possess inherent order whereas sets do not.

Our target output is a sequence of integer pointers $Y = \langle y_1, \dots, y_n : y_i \in \mathbb{Z}^+ \rangle$, satisfying the conditions: $\forall i, j : y_i \neq y_j$ and $\forall i, 1 \leq y_i \leq n$, signifying that no element from the input can take multiple positions in the output sequence and every element must be assigned a position in the final reordering. Our objective is to define a differentiable function $f : X \rightarrow Y$ that meets these criteria, with some parameters θ , lending itself to gradient-based training.

The simplest effective formulation of such an attention-based ordering function requires the following learnable parameters: two square weights matrices $W_1, W_2 \in \mathbb{R}^{d \times d}$ as well as a single vector $\mathbf{v} \in \mathbb{R}^{d \times 1}$. Additionally, given the sequential nature in which the ordering is generated, we also have access to a decoder state \mathbf{s}_t at each step t , for which we calculate the attention-based probability distribution over input elements. The array of encoded set elements X is processed in the following way:

$$z_t^i = \mathbf{x}_i W_1 + \mathbf{s}_t W_2 \quad \text{for } i \in (1, \dots, n) \quad (19)$$

$$a_t^i = \frac{\exp(\mathbf{v}^T \sigma(z_t^i))}{\sum_j^n \exp(\mathbf{v}^T \sigma(z_t^j))} \quad \text{for } i \in (1, \dots, n) \quad (20)$$

$$y_t = \operatorname{argmax}(\mathbf{a}_t) \quad \text{for } t \in (1, \dots, n) \quad (21)$$

where σ is a nonlinear activation function, commonly the hyperbolic tangent (Logeswaran et al., 2018). The vector $\mathbf{z}_t = \langle z_t^1, \dots, z_t^n \rangle$ is a representation of all input elements, adjusted by the representation of the generated output sequence at step t , through additive attention (Bahdanau et al., 2015). Finally, y_t is an integer pointer to the element of X which received the highest attention value within the attention vector $\mathbf{a}_t = \langle a_t^1, \dots, a_t^n \rangle$ at the current step t of the sequentially predicted output sequence.

Most methods obtain the embedding of the entire input X through more complex encoding methods, such as stacked bidirectional LSTMs, as seen in the work of Vinyals et al. (2015) and Vinyals et al. (2016). This representation retains the information about the original arbitrary order of the elements and should be dependent on the entire set. In the above example, the length t of the output sequence of pointers Y is the same as the cardinality of X' , which is the case in many combinatorial optimization problems. However, in some structured prediction tasks the optimal length of the output is also subject to the learning process, allowing for pointing to the same element \mathbf{x}_i multiple times or not at all, as is the case in catalog design (Carlson-Skalak et al., 1998).

It should be noted that in the above formulation nothing is explicitly preventing the model from pointing to the same element \mathbf{x}_i at multiple positions of the output sequence. In practice the iterative learning process can largely prevent this from occurring, given appropriate training data. One method to explicitly prohibit repetition is progressive masking, as described in section 3.2.1.

In practice, beam search is commonly used during inference to increase the probability that the most optimal sequence is predicted, as evidenced in the work of Bahdanau et al. (2017) and Kool et al. (2019). Beam search employs a heuristic search algorithm to expand a limited number of most promising vertices of the graph defined by the attention vectors over a predefined number of output steps. It tracks a small number of potential partial solutions, at the controlled cost of memory and computation. However, the exact reason why a small beam number results in qualitatively better predictions has recently been called into question (Meister et al., 2020). For a more detailed overview in the context of sequence prediction, see Wiseman and Rush (2016).

4.2.2 METHODS USING ATTENTION-BASED POINTING

Many neural network model architectures across distant fields of application include variations on the basic attention-based pointing mechanism. These include specifically set-to-sequence methods, such as the Pointer Networks discussed in section 3.2.2 and the RPW model detailed in section 3.3.2, as well as elements of complex models in reinforcement learning applied to competitive real-time strategy challenges (Vinyals et al., 2019).

Additionally, attention-based pointing found usage in identifying entailment between documents (Rocktäschel et al., 2016), abstractive text summarization (See et al., 2017), rare and out-of-vocabulary word prediction (Merity et al., 2016) as well as describing multimedia content (Cho et al., 2015). An interesting augmentation of attention-based pointing has been experimented with in the context of generating structured queries (in SQL) from natural language sequences (Zhong et al., 2018). However, the improvement stems from augmenting the input with SQL keywords to limit the output space, not from an essential adjustment to the underlying attention-based pointing mechanism.

An application of attention-based pointing to generate solutions to another classic combinatorial optimization challenge, the Vehicle Routing Problem, was proposed by Kool et al. (2019). The resulting architecture is an encoder-decoder *Graph Attention Network* (Veličković et al., 2018), employing multiheaded attention in the encoder and node masking in the decoder, which uses the embeddings of all the nodes and the entire graph at each step t to point to the next node to be visited. Unlike the previously discussed methods, this model is trained using a gradient estimator from the field of reinforcement learning, first proposed by Williams (1992).

4.3 Permutation Matrices

A permutation matrix is a square, binary matrix P having exactly one entry $p_j^i = 1$ in each row i and each column j (Stuart & Weaver, 1991). All other entries are equal to 0. P 's dimensions are defined by the length of the desired output sequence, most commonly equal to the length of the input. A permutation matrix is unimodal in that each of its rows has a unique, highest value. It is also doubly-stochastic, since it consists entirely of nonnegative numbers, with each row and column summing to 1.

The key property of a permutation matrix is that given an arbitrarily ordered sequence, we can left-multiply it by a permutation matrix to obtain a reordered sequence. For example, given an arbitrarily shuffled input sequence $\mathbf{x} = \langle x_1, x_5, x_3, x_4, x_2 \rangle$, which we want to permute in such a way as to restore the i -indexed ascending order to obtain $\mathbf{x}^\pi = \langle x_1, x_2, x_3, x_4, x_5 \rangle$, we can predict the permutation matrix visualized below. Our target is the permutation π , given in two-line notation:

$$\pi = \begin{pmatrix} x_1 & x_5 & x_3 & x_4 & x_2 \\ 1 & 5 & 3 & 4 & 2 \end{pmatrix} \quad (22)$$

We can apply this permutation to the transposed input \mathbf{x}^\top via left-multiplication by the target permutation matrix P :

$$\mathbf{x}^\pi = P\mathbf{x}^\top = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_5 \\ x_3 \\ x_4 \\ x_2 \end{bmatrix} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \end{bmatrix} \quad (23)$$

This gives us a way to represent the target permutation in the form of a matrix of numbers between 0 and 1, marking an important step towards enabling the application of machine learning methods.

4.3.1 MAKING PERMUTATION MATRICES LEARNABLE

In order to use a permutation matrix as the target output of a machine learning model that can be trained through some form of gradient-based iterative optimization, we must introduce certain relaxations of the concept. Otherwise, we are left with the equivalent of a sorting operator which is non-differentiable (Grover et al., 2019).

Traditionally, sorting operations result in either a permutation $\pi = \langle 1, 5, 3, 4, 2 \rangle$ (from the previous example) or a vector of reordered elements $\mathbf{x}^\pi = \langle x_1, x_2, x_3, x_4, x_5 \rangle$. Much like the permutation matrix, the former is non-differentiable with respect to the input due to being integer-valued and the latter due to being piecewise linear (Cuturi et al., 2019).

The most direct way to obtain a differentiable relaxation P' of the permutation matrix P is to map the input \mathbf{x} to a continuous codomain, as opposed to the original discrete one. Thus, we need a relaxation such that $P' \in \mathbb{R}^{n \times n}$, lending itself to gradient-based optimization. Additionally, an efficient projection from the continuous codomain back to the discrete one must exist to allow for the use of loss functions and evaluation metrics. This can be achieved by applying the $\operatorname{argmax}()$ function per row of P' to find the position in which the unique 1 would have been located in the actual permutation matrix P .

The resulting relaxation must retain the property of row-stochasticity, such that:

$$\forall i, j \in \{1, \dots, n\} : p'_{ij} \geq 0 \tag{24}$$

$$\forall i \in \{1, \dots, n\} : \sum_{j=1}^n p'_{ij} = 1 \tag{25}$$

and of unimodality, such that a vector \mathbf{y} is both obtainable from the relaxation matrix P' through row-wise $\operatorname{argmax}()$ and a valid permutation of the input:

$$\forall i \in \{1, \dots, n\} : y_i = \operatorname{argmax}(\mathbf{p}'_i) \tag{26}$$

There are many possible methods of obtaining this relaxation from the input after it has been transformed by a chosen neural network architecture, such as adding elementwise Gumbel perturbations (Mena et al., 2018), applying the Sinkhorn operator to directly sample matrices near the Birkhoff polytope (Linderman et al., 2018), which is the convex hull whose points are doubly-stochastic matrices (Emami & Ranka, 2018), or through the application of a $\operatorname{softmax}()$ operator on a derived matrix of absolute pairwise distances between the individual input elements (Grover et al., 2019).

Depending on the specifics of the task at hand, the target matrix can be predicted in a single pass (if the length of the input and output is a constant, known prior to inference) or sequentially, row by row (Nishida & Nakayama, 2017). An interesting method employing permutation matrices has been proposed by Zhang et al. (2019), in which a trainable, pairwise *ordering cost* function is used to produce an anti-symmetric matrix C , whose entry c_{ij} represents the cost of placing the i -th element before the j -th. This function is represented as a neural network, which is then used to continuously adjust the learned permutation matrix. This is referred to as a *Permutation-Optimisation* module, and has been demonstrated to perform well on number sorting, re-assembling image mosaics and visual question answering, with one limiting feature of entailing cubic time complexity.

4.4 Listwise Ranking

Ranking methods determine the predicted order of elements by assigning a score to each element and then sorting them according to these scores. In listwise ranking, a score is calculated based on a list of available elements and a *query*, for which a specific order is to be predicted. The terminology stems from applications in information retrieval, where the task is to rank the available *documents*

(such as web pages) in order of relevance for a given query (e.g. a search term). For an overview of neural ranking methods, see Mitra et al. (2018).

Listwise ranking is distinguished from point- and pair-wise ranking. Pointwise methods reduce the ranking problem to regression, in that the relevance score for an element is obtained only from its own representation and the query. Pairwise approaches reduce it to binary classification (Lei et al., 2017). Given a query and a pair of elements, they predict which element is the more relevant of the two. In listwise ranking, the prediction is performed on a list of objects.

Application of listwise ranking methods to set-to-sequence problems requires certain adjustments. Most importantly, in such ordering and structure prediction challenges we are not given a specific query, for which an ordering is to be generated. However, we are able to obtain a permutation invariant set representation through the set encoding methods detailed in section 3. We can use this learned embedding to fill the role that a query performs in traditional learn-to-rank methods. Intuitively, the relative rank of each element in the output sequence should depend on the entire available set.

This marks an important departure from the assumption that two elements have a canonical relative order, which should remain unchanged regardless of what other elements are present in the input set. In practice, higher-order interactions between available elements can entirely change the target sequence. In order to learn such properties, a ranking method can utilize the permutation invariant set representation as the query and predict the relevance scores of all available elements in one go, in the listwise manner described in the following subsection.

4.4.1 DETAILS OF LISTWISE RANKING

In order to illustrate the basic underlying mechanisms in listwise ranking methods, this subsection investigates the first method that formulated order prediction on a list of objects, namely the *ListNet* model, proposed by (Cao et al., 2007). From this point on, whenever we refer to a query, we are referring to the permutation invariant representation of the entire set obtained via the previously outlined set encoding methods. This representation can be learned in parallel with the weights of the ranking the model.

Once more, assume we are given an arbitrarily ordered input sequence $X = \langle \mathbf{x}_1, \dots, \mathbf{x}_n \rangle$ such that $X \in \mathbb{R}^{n \times d}$. This sequence is one of $n!$ possible permutations of the corresponding set $X' = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$. The objective is to transform the given X vector into the best possible permutation of X' , according to some preference. This target permutation is $\pi = \langle \pi(1), \dots, \pi(n) \rangle$, such that $\pi(i)$ is the object at position i in the permutation. The set of all possible permutations of length n is denoted as Ω_n . The target permutation π is represented by a vector of scores $\mathbf{y}^\pi = \langle y_{\pi(1)}, \dots, y_{\pi(n)} \rangle$, where $y_{\pi(i)}$ is the predicted relevance score for the element \mathbf{x}_i . All elements of the original input X get assigned a score by the learned neural network, in a listwise manner.

A naive approach would be to try to obtain the probability $P_y(\pi)$ of each of the $n!$ possible permutations from Ω_n , given a set of scores \mathbf{y} :

$$P_y(\pi) = \prod_{i=1}^n \frac{\phi(y_{\pi(i)})}{\sum_{k=i}^n \phi(y_{\pi(k)})} \quad (27)$$

where $\phi()$ is any increasing and strictly positive function applied to the scores and $y_{\pi(i)}$ is the relevance score of element i in the vector \mathbf{y}_π corresponding to some permutation π (Cao et al., 2007).

However, this is a computationally inefficient approach, therefore we instead calculate the *top one probability* for each element \mathbf{x}_i . This probability is equal to the sum of the permutation probabilities of every permutation where the i -th element was ranked first. For the exact proof of this equivalence, see the appendix in Cao et al. (2007). The crucial observation is that we do not need to calculate all the permutation probabilities to obtain the correct top one probability $P_{top}(\mathbf{x}_i)$ of each element, given a list of scores \mathbf{y}_π :

$$P_{top}(\mathbf{x}_i) = \frac{\exp(y_{\pi(i)})}{\sum_{k=1}^n \exp(y_{\pi(k)})} \quad (28)$$

In effect, the probability distribution over permutations is obtained by applying the softmax() function to the predicted relevance scores. Since this is a supervised learning framework, our model has access to the ground truth distribution for each example list. This allows for the use of loss functions that compare two distributions, such as the Kullback-Leibler (KL) divergence, also referred to as relative entropy (Liu & Shum, 2003).

In order to obtain the vector of relevance scores \mathbf{y}_π , the query vector $\mathbf{q} \in \mathbb{R}^m$ is concatenated with a fixed-length, learned embedding of each element of the input set \mathbf{x}_i , resulting in the final input $I \in \mathbb{R}^{n \times (m+d)}$, per single example list. In the original ListNet model, this input is transformed through stacked, fully-connected layers, of which the last one has a single unit outputting the relevance score per element.

As such, this transforms each concatenation of the query and the embedded element in an independent and identical manner. This is practical for ranking tasks, where two items always have the same relative target order, given a query. However, in more complex structure prediction tasks, common in set-to-sequence challenges, relying entirely on the set encoding \mathbf{q} and the listwise loss to identify interactions between set elements limits the ability to approximate complex functions, in which higher-order interactions between elements have an impact on the final order (Lan et al., 2009). This limitations of ListNet is partially addressed by later methods, discussed in section 4.4.2.

Additionally, a limitation of the ranking approach is that by representing the target output sequence through relevance scores per input element, we cannot learn to output sequences of any length other than n . This also precludes the use of listwise ranking in challenges where the output sequence is of the same length as the number of input elements n , but repetition and exclusion of elements is required. For a more detailed look into the underlying theoretical aspects of listwise ranking, see Xia et al. (2008).

4.4.2 METHODS USING LISTWISE RANKING

Multiple listwise ranking methods in machine learning have been developed since the original ListNet model. A notable framework that addresses the challenge of learning pairwise interactions between elements came in the form of *BoltzRank*, in which the rank probabilities are sampled from a Boltzmann distribution, employing an energy function that depends on a score influenced by both individual and pairwise potentials (Volkovs & Zemel, 2009).

Another approach of interest is the *FATE* framework proposed by Pfannschmidt et al. (2018). The authors identify the problem of predicting a relevance score for an element and the query with only the loss function carrying the signal regarding the context of other available elements. In order to address this, they effectively employ a permutation invariant set-encoder (Zaheer et al., 2017), whose output is concatenated to the learned representation of each element, in a variation of the basic method described in section 4.4.1.

Ai et al. (2018) propose a complex architecture combining two-stage ranking, sequential recurrent neural networks and an attention-based loss function. The proposed *Deep Listwise Context Model* (DLCM) sequentially encodes the most relevant results using the corresponding feature vectors, trains an additional, local context model and employs it to re-rank the best k results.

Finally, Pang et al. (2020) employ the Set Transformer to obtain element representations that encode cross-document interactions and return a permutation invariant ranking by sorting the permutation equivariant relevance scores per each document.

5. Discussion

In this section, the key challenges and contexts of application for set-to-sequence deep learning methods are discussed. Additionally, separate sections are devoted to the progress in set encoding and permutation learning, along with current limitations and proposed directions for future research.

5.1 Key Challenges

Core challenges of set-to-sequence include representing sets of varying sizes in a permutation invariant way, encoding pairwise and higher-order interactions between set elements and keeping the computational complexity moderate in the cardinality of the input set. This last requirement stems from the fact that in many important set-to-sequence problems the size of the input set can be quite large, particularly in computer vision point-cloud based tasks (Ge et al., 2018). This has been addressed by various proposed set-encoding or set-pooling methods discussed in section 3, such as RNN-based Pointer Networks (Vinyals et al., 2015) and Read-Process-and-Write models (Vinyals et al., 2016) as well as through fully connected methods such as DeepSets (Zaheer et al., 2017) and the Set Transformer (Lee et al., 2019), among others.

Another set of key challenges in set-to-sequence stems from the many possible ways of formalizing the process of outputting a permutation sequence. Potential methods include learn-to-rank approaches, permutation matrices and attention-based methods discussed in section 4. Additionally, it is not always trivial to combine the output of the specific set-encoding technique with the expected input of these permutation methods in a way that results in good model performance. Finally, we expect not just the set-encoding module but the entire set-to-sequence model to be permutation invariant and always give the optimal output regardless of how the array representing the set is reordered at input.

5.2 Contexts of Application

There are many contexts where the input data does not have an inherent ordering and the number of input elements varies (i.e. *set-input* problems) and possibly even more where the input elements are not unique, instead repeating a meaningful number of times, in which case the task presents an *mset*- or *multiset-input* problem. Additionally, the set elements may be allowed to reoccur multiple times in the predicted sequence or be excluded from it entirely, forming a *complex permutation*. These elements can be both continuous (e.g. word embeddings) or discrete (coming from an index-based dictionary). An example of such a problem with a strong industrial application is the question of how to order available product offers into a displayable catalogue that will keep the reader engaged and eventually inspire them to make a relevant purchase (Liao & Chen, 2004).

Similar challenges are faced by the experts who order news articles into a coherent publication, book authors composing chapters into a novel and by engineers tackling the challenge of *catalog design*, in which a configuration is created by assembling off-the-shelf components into a functional system (Carlson-Skalak et al., 1998). Set-to-sequence methods can also benefit architects in their search for the best configurations of the design space, taking into consideration structural efficiency, daylight availability and other aspects of building performance (Brown & Mueller, 2017). Whenever the elements come from a pre-existing set and the output is structured as a complex permutation, we are facing a set-to-sequence challenge.

5.3 Progress in Set Encoding

As machine learning methods become more and more widely used, the range of input and output data structures that these methods are applied to becomes larger. Mirroring the development and consecutive growth in popularity of methods created specifically to handle certain types of input data, such as convolutional neural networks for images and recurrent models for sequences, one of the

most significant developments in set-to-sequence has been the work towards obtaining permutation invariant representations of sets of varying lengths.

Particularly, this representation can now be obtained in a way that is capable of explicitly encoding not just pairwise but also higher-order interactions between input elements. Intuitively these improvements mirror the way we as humans process input sets, in that we group the elements into meaningful collections, allowing for enough flexibility that this grouping may change entirely upon the introduction of a new element.

5.4 Progress in Permutation Learning

Comparative progress has been made regarding the ways in which we formalize the permutation tasks at hand. Given that the simplest form of permutation (reordering) is sorting, the early methods formalized this challenge as a ranking problem, with pointwise, pairwise and eventually listwise loss being used to train the model. These methods suffer from the assumption that any two elements have a canonical pairwise ranking, regardless of the features of the other elements in the input set.

Another alternative emerged in the form of self-attention applied in such a way as to output softmaxed pointers to the elements of the input set at each step in the output sequence, with the additional use of beam search during inference (Vinyals et al., 2015). The third dominant formalization comes in the form of learning to output a permutation matrix. This allows for the original input to then be matrix-multiplied by the row-stochastic permutation matrix, resulting in the desired reordering (Nishida & Nakayama, 2017), (Emami & Ranka, 2018). It is a sign of the complexity of the field that no clear preferred formalization of its core challenge has emerged, with ranking (for example) still finding useful application in active research (Kumar et al., 2020).

5.5 Limitations

A crucial limitation of many of the cutting edge set-to-sequence models, such as the Set Transformer (Lee et al., 2019) and the permutation mechanism in Pointer Networks, is their reliance on self-attention. Whilst Transformer-based methods that rely solely on multiheaded self-attention have seen remarkable success, even in applications beyond a fixed-length context (Dai et al., 2019), their ability to process hierarchical structure has hard limits. Specifically, purely self-attention architectures are entirely dependent on the number of multi-attention heads and layers growing alongside the size of the input to retain the ability to model recursion, finite-state languages and other hierarchical aspects of our data (Hahn, 2020).

In order to overcome this limitation, efforts have been made to combine the benefits of sequential computation, inherent to recurrent neural networks, with the advantages of parallel computation and the global receptive field of the Transformer. This method is referred to as the *Universal Transformer* (Dehghani et al., 2019). In addition, it also includes the *Adaptive Computation Time* mechanism proposed by Graves (2016), which enables the model to dynamically learn how many computational steps to perform per the features of the input sequence. However, these advances have not yet been translated to the domain of set-input challenges.

5.6 Directions for Further Research

There are many possible directions for further investigations, pertaining to both the area of encoding sets and proposing novel permutation learning methods. As hinted in section 4.1, the question of how best to utilize the fixed length set representation within the internals of the ordering module remains open. However, we believe the areas of learning complex permutations (5.6.1) and differentiable loss functions (5.6.2) deserve separate attention, given in the following subsections.

5.6.1 COMPLEX PERMUTATIONS

An important and natural extension of the research within set-to-sequence is to apply it to challenges where the output sequence allows for repetition and exclusion of input set elements, thus going beyond traditional permutation learning (Diallo et al., 2020). As such, these *complex permutations* present an additional challenge of dynamically predicting the optimal sequence length without sacrificing the second of the two aforementioned prerequisites for ML methods on sets, namely the requirement that the same model must be able to process finite input sets of any cardinality. It is not immediately clear how to achieve both properties, with multiple promising approaches gaining prominence, ranging from the addition of a confidence loss for long time-series prediction (Harmon & Klabjan, 2019), to the aforementioned adaptive computation time (Wu et al., 2020). For an overview of dynamic neural network methods in general, see Han et al. (2021).

Important tasks involving a complex target permutation include predicting a configuration representing the assembly of off-the-shelf components into a functional system (Carlson-Skalak et al., 1998) and the composition of product offers into rendered catalogues (Liao & Chen, 2004). Additionally, an interesting area for further research would be to extend these complex permutation sequences to grids and lattices, as suggested by Zhang et al. (2019), or even to graphs, expanding on the work of Serviansky et al. (2020).

5.6.2 DIFFERENTIABLE LOSS FUNCTIONS

Another important area for further work is centered around the problem of framing the set-to-sequence challenges in such a way as to enable the use of differentiable loss functions. A naive but often practically effective approach is to frame the problem as categorization and use a cross-entropy loss, as discussed in Engilberge et al. (2019). However, it precludes meaningful distinction between pointing to an incorrect element that is very similar to the correct one and pointing to an entirely different incorrect element.

Alternatively, if a ranking framework is applied, where a score is generated for each element and subsequently used to sort all elements into a new permutation, we gain access to well documented listwise losses, such as the ones successfully employed in the ListNet or *ListMLE* (Kumar et al., 2020) frameworks. Many metrics that would lend themselves to ordering challenges do not have defined derivatives for their entire domain. The development and testing of their smooth approximations is of great potential value, as seen in the works of Rolínek et al. (2020) and Blondel et al. (2020).

6. Conclusion

Set-to-sequence is currently established as a family of methods with an exceptionally wide range of applications. At its essence, it is a combination of three areas seeing a lot of attention within the machine learning and deep learning research communities, namely set-input problems, combinatorial optimization and structured prediction. As such, a number of methods that can further our understanding of this field originates from other areas of interest, sometimes without seeing immediate application to set-to-sequence challenges. Progress owes to advances in model architectures, parallel computing, hyperparameter optimization as well as the overall growing interest in applying machine learning solutions to a wider and wider gamut of industrial challenges.

We have presented an overview of set-to-sequence methods from the fields of machine learning and deep learning. The initial sections of this paper introduced the reader to the key concepts in relation to machine learning on sets, most notably the properties of permutation invariance, permutation equivariance and the requirement of handling inputs of varying lengths. In relation to this, in section 3, the reader is introduced to a number of selected set-encoding model architectures, including a significantly detailed look at their underlying mathematical transformations. To facilitate comprehensive understanding, we compared these and other related methods through several summary tables presented in section 3.1.

Additionally, a survey of potential ordering methods has been provided in section 4. This included the three primary ways of formalizing the output permutation in the set-to-sequence setting, namely listwise ranking, relaxations of permutation matrices and attention-based pointing. Once a permutation invariant set representation is obtained through one of the aforementioned set-encoding models, the ordering methods’ calculations can be conditioned on this embedding to output a sequence of elements from the original input set, in one fully trainable deep learning framework.

Acknowledgements

This work was partly supported by an Innovation Fund Denmark research grant (number 9065-00017B) and by Tjek A/S. The authors would like to acknowledge Rasmus Pagh’s assistance in the conceptualization of different set-to-sequence settings and general feedback on the early drafts of this paper.

References

- Achlioptas, P., Diamanti, O., Mitliagkas, I., & Guibas, L. (2018). Learning representations and generative models for 3D point clouds. In *6th International Conference on Learning Representations, ICLR 2018 - Workshop Track Proceedings*.
- Ai, Q., Bi, K., Guo, J., & Croft, W. B. (2018). Learning a deep listwise context model for ranking refinement. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pp. 135–144.
- Altae-Tran, H., Ramsundar, B., Pappu, A. S., & Pande, V. (2017). Low data drug discovery with one-shot learning. *ACS central science*, 3(4), 283–293.
- Arora, S. (1996). Polynomial time approximation schemes for euclidean tsp and other geometric problems. In *Proceedings of 37th Conference on Foundations of Computer Science*, pp. 2–11. IEEE.
- Ba, J. L., Kiros, J. R., & Hinton, G. E. (2016). Layer normalization. *arXiv preprint arXiv:1607.06450, ArXiv*.
- Bahdanau, D., Brakel, P., Xu, K., Goyal, A., Lowe, R., Pineau, J., Courville, A., & Bengio, Y. (2017). An Actor-Critic Algorithm for Sequence Prediction. In *ICLR*, pp. 1–17.
- Bahdanau, D., Cho, K., & Bengio, Y. (2015). Neural machine translation by jointly learning to align and translate. *CoRR*, abs/1409.0473.
- Bello, I., Kulkarni, S., Jain, S., Boutilier, C., Chi, E., Eban, E., Luo, X., Mackey, A., & Meshi, O. (2018). Seq2slate: Re-ranking and slate optimization with rnns. In *Proceedings of the Workshop on Negative Dependence in Machine Learning at the 36th International Conference on Machine Learning*,.
- Bello, I., Pham, H., Le, Q. V., Norouzi, M., & Bengio, S. (2016). Neural combinatorial optimization with reinforcement learning. *arXiv preprint arXiv:1611.09940, ArXiv*.
- Bengio, Y., Lodi, A., & Prouvost, A. (2020). Machine learning for combinatorial optimization: a methodological tour d’horizon. *European Journal of Operational Research*, 1.
- Blondel, M., Teboul, O., Berthet, Q., & Djolonga, J. (2020). Fast differentiable sorting and ranking. In *Proceedings of the 37th International Conference on Machine Learning, PMLR 119, 2020*.
- Brown, N., & Mueller, C. (2017). Designing with data: Moving beyond the design space catalog. *Disciplines and Disruption - Proceedings Catalog of the 37th Annual Conference of the Association for Computer Aided Design in Architecture, ACADIA 2017, 1*, 154–163.

- Brown, T., Mann, B., Ryder, N., & Subbiah, M. (2020). Language models are few-shot learners. *ArXiv, abs/2005.14165*.
- Cao, Z., Qin, T., Liu, T.-Y., Tsai, M.-F., & Li, H. (2007). Learning to rank: from pairwise approach to listwise approach. In *Proceedings of the 24th international conference on Machine learning*, pp. 129–136.
- Carlson-Skalak, S., White, M. D., & Teng, Y. (1998). Using an evolutionary algorithm for catalog design. *Research in Engineering Design, 10*, 63–83.
- Chaudhari, S., Polatkan, G., Ramanath, R., & Mithal, V. (2019). An attentive survey of attention models. *arXiv preprint arXiv:1904.02874, ArXiv*.
- Cheng, J., Dong, L., & Lapata, M. (2016). Long short-term memory-networks for machine reading. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 551–561, Austin, Texas. Association for Computational Linguistics.
- Cho, K., Courville, A., & Bengio, Y. (2015). Describing multimedia content using attention-based encoder-decoder networks. *IEEE Transactions on Multimedia, 17*(11), 1875–1886.
- Cohen, W. W., Schapire, R. E., & Singer, Y. (1998). Learning to order things. In *Advances in Neural Information Processing Systems*, pp. 451–457.
- Cui, B., Li, Y., Chen, M., & Zhang, Z. (2018). Deep attentive sentence ordering network. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, EMNLP 2018*, pp. 4340–4349.
- Cui, G., Yu, W., Li, X., Zeng, Z., & Gu, B. (2019). Machine-Learning-Driven Matrix Ordering for Power Grid Analysis. In *Proceedings of the 2019 Design, Automation and Test in Europe Conference and Exhibition, DATE 2019*, pp. 984–987. EDAA.
- Cuturi, M., Teboul, O., & Vert, J.-P. (2019). Differentiable ranking and sorting using optimal transport. In *Advances in Neural Information Processing Systems*, pp. 6861–6871.
- Dai, H., Khalil, E. B., Zhang, Y., Dilkina, B., & Song, L. (2017). Learning combinatorial optimization algorithms over graphs. *Advances in Neural Information Processing Systems, 2017-December(Nips)*, 6349–6359.
- Dai, Z., Yang, Z., Yang, Y., Carbonell, J., Le, Q., & Salakhutdinov, R. (2019). Transformer-XL: Attentive language models beyond a fixed-length context. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 2978–2988, Florence, Italy. Association for Computational Linguistics.
- De Cao, N., & Kipf, T. (2018). MolGAN: An implicit generative model for small molecular graphs. *ICML Deep Generative Models Workshop, 1*.
- Dehghani, M., Gouws, S., Vinyals, O., Uszkoreit, J., & Kaiser, L. (2019). Universal transformers. In *7th International Conference on Learning Representations, ICLR 2019*, pp. 1–23.
- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Diallo, A., Zopf, M., & Fürnkranz, J. (2020). Permutation learning via lehmer codes. *Frontiers in Artificial Intelligence and Applications, 325*, 1095–1102.
- Edwards, H., & Storkey, A. (2017). Towards a Neural Statistician. In *5th International Conference on Learning Representations, ICLR 2017 - Conference Track Proceedings*, pp. 1–13.
- Emami, P., & Ranka, S. (2018). Learning permutations with sinkhorn policy gradient. *ArXiv, abs/1805.07010*.

- Engilberge, M., Chevallier, L., Pérez, P., & Cord, M. (2019). Sodeep: A sorting deep net to learn ranking loss surrogates. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Ge, L., Cai, Y., Weng, J., & Yuan, J. (2018). Hand pointnet: 3d hand pose estimation using point sets. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 8417–8426.
- Graves, A. (2016). Adaptive Computation Time for Recurrent Neural Networks. *Preprint arXiv:1603.08983, ArXiv*.
- Graves, A., Wayne, G., & Danihelka, I. (2014). Neural turing machines. *Preprint arXiv:1410.5401, ArXiv*.
- Grinman, A. (2015). The hungarian algorithm for weighted bipartite graphs. *Massachusetts Institute of Technology, MIT*.
- Grover, A., Wang, E., Zweig, A., & Ermon, S. (2019). Stochastic optimization of sorting networks via continuous relaxations. *Preprint arXiv:1903.08850, ArXiv*.
- Hahn, M. (2020). Theoretical Limitations of Self-Attention in Neural Sequence Models. *Transactions of the Association for Computational Linguistics*, 8.
- Halmos, P. R. (2017). *Naive set theory*. Courier Dover Publications.
- Hamilton, W. L., Ying, Z., & Leskovec, J. (2017). Inductive representation learning on large graphs. In *NIPS*.
- Han, Y., Huang, G., Song, S., Yang, L., Wang, H., & Wang, Y. (2021). Dynamic neural networks: A survey. *Preprint arXiv:2102.04906, ArXiv*.
- Harmon, M., & Klabjan, D. (2019). Dynamic prediction length for time series with sequence to sequence networks. *arXiv*, 1(2014), 1–23.
- Hayden, S., Zermelo, E., Fraenkel, A. A., & Kennison, J. F. (1968). *Zermelo-Fraenkel set theory*. CE Merrill.
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778.
- Hernández-García, A., König, P., & Kietzmann, T. (2019). Learning robust visual representations using data augmentation invariance. *ArXiv*, *abs/1906.04547*.
- Hu, D. (2019). An introductory survey on attention mechanisms in nlp problems. In *Proceedings of SAI Intelligent Systems Conference*, pp. 432–448. Springer.
- Huang, G., Liu, S., Van der Maaten, L., & Weinberger, K. Q. (2018a). Condensenet: An efficient densenet using learned group convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2752–2761.
- Huang, P. S., Wang, C., Singh, R., Yih, W. T., & He, X. (2018b). Natural language to structured query generation via meta-learning. *NAACL HLT 2018 - 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies - Proceedings of the Conference*, 2, 732–738.
- Ibrahim, A., R, A., & Ishaya, J. (2019). The Capacitated Vehicle Routing Problem. *International Journal of Research - Granthaalayah*, 3(2), 30–33.
- Ionescu, C., Vantzos, O., & Sminchisescu, C. (2015). Matrix backpropagation for deep networks with structured layers. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2965–2973.
- Iqbal, S., & Sha, F. (2019). Actor-attention-critic for multi-agent reinforcement learning. In *International Conference on Machine Learning*, pp. 2961–2970. PMLR.

- Ishaya, J., Ibrahim, A., & Lo, N. (2019). A comparative analysis of the travelling salesman problem: Exact and machine learning techniques. *Open Journal of Discrete Applied mathematics*, 2(3), 23–37.
- Jain, S., & Wallace, B. C. (2019). Attention is not explanation. *Preprint arXiv:1902.10186, ArXiv*.
- Jaiswal, A., Wu, R. Y., Abd-Almageed, W., & Natarajan, P. (2018). Unsupervised adversarial invariance. In *Advances in Neural Information Processing Systems*, pp. 5092–5102.
- Johnson, J., Hariharan, B., Van Der Maaten, L., Fei-Fei, L., Lawrence Zitnick, C., & Girshick, R. (2017). Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2901–2910.
- Jung, I.-S., Berges, M., Garrett, J. H., & Poczos, B. (2015). Exploration and evaluation of ar, mpca and kl anomaly detection techniques to embankment dam piezometer data. *Adv. Eng. Inform.*, 29(4), 902–917.
- Karch, T., Colas, C., Teodorescu, L., Moulin-Frier, C., & Oudeyer, P.-Y. (2020). Deep sets for generalization in rl. *Preprint arXiv:2003.09443, ArXiv*.
- Karlin, A. R., Klein, N., & Gharan, S. O. (2020). A (slightly) improved approximation algorithm for metric tsp. *ArXiv, abs/2007.01409*.
- Kayhan, O. S., & Gemert, J. C. v. (2020). On translation invariance in cnns: Convolutional layers can exploit absolute spatial location. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14274–14285.
- Kersting, K., Kriege, N. M., Morris, C., Mutzel, P., & Neumann, M. (2016). Benchmark data sets for graph kernels. *URL <http://graphkernels.cs.tu-dortmund.de>, 29*.
- Kool, W., Van Hoof, H., & Welling, M. (2019). Attention, learn to solve routing problems!. In *7th International Conference on Learning Representations, ICLR 2019*, pp. 1–25.
- Kosioerek, A. R., Kim, H., & Rezende, D. J. (2020). Conditional Set Generation with Transformers. In *Workshop on Object-Oriented Learning at ICML 2020*.
- Kumar, P., Brahma, D., Karnick, H., & Rai, P. (2020). Deep Attentive Ranking Networks for Learning to Order Sentences. In *AAAI 2020*.
- Lake, B. M., Salakhutdinov, R., & Tenenbaum, J. B. (2019). The omniglot challenge: a 3-year progress report. *Current Opinion in Behavioral Sciences*, 29, 97–104.
- Lan, Y., Liu, T.-Y., Ma, Z., & Li, H. (2009). Generalization analysis of listwise learning-to-rank algorithms. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pp. 577–584.
- Lee, J., Lee, Y., Kim, J., Kosioerek, A. R., Choi, S., & Teh, Y. W. (2019). Set Transformer: A Framework for Attention-based Permutation-Invariant Neural Networks. *Proceedings of the 36th International Conference on Machine Learning*, 36.
- Lei, Y., Li, W., Lu, Z., & Zhao, M. (2017). Alternating pointwise-pairwise learning for personalized item ranking. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, pp. 2155–2158.
- Liao, S. H., & Chen, Y. J. (2004). Mining customer knowledge for electronic catalog marketing. *Expert Systems with Applications*, 27(4), 521–532.
- Lin, T.-Y., Maji, S., & Koniusz, P. (2018). Second-order democratic aggregation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 620–636.
- Linderman, S., Mena, G., Cooper, H., Paninski, L., & Cunningham, J. (2018). Reparameterizing the birkhoff polytope for variational permutation inference. In *International Conference on Artificial Intelligence and Statistics*, pp. 1618–1627. PMLR.

- Ling, J., Jones, R., & Templeton, J. (2016). Machine learning strategies for systems with invariance properties. *Journal of Computational Physics*, 318, 22–35.
- Liu, C., & Shum, H.-Y. (2003). Kullback-leibler boosting. In *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings.*, Vol. 1, pp. I–I. IEEE.
- Liu, X., Weijer, J. V. D., & Bagdanov, A. D. (2019). Exploiting Unlabeled Data in CNNs by Self-Supervised Learning to Rank. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(8), 1862–1878.
- Liu, Z., Luo, P., Wang, X., & Tang, X. (2015). Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*.
- Locatello, F., Weissenborn, D., Unterthiner, T., Mahendran, A., Heigold, G., Uszkoreit, J., Dosovitskiy, A., & Kipf, T. (2020). Object-Centric Learning with Slot Attention. *arXiv*, 1(NeurIPS), 1–27.
- Logeswaran, L., Lee, H., & Radev, D. (2018). Sentence ordering and coherence modeling using recurrent neural networks. In *32nd AAAI Conference on Artificial Intelligence, AAAI 2018*, pp. 5285–5292.
- Loosli, G., Canu, S., & Bottou, L. (2007). Training invariant support vector machines using selective sampling. *Large scale kernel machines*, 2.
- Ma, S., Deng, Z.-H., & Yang, Y. (2016). An unsupervised multi-document summarization framework based on neural document model. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pp. 1514–1523.
- Mani, K., Verma, I., Meisheri, H., & Dey, L. (2018). Multi-document summarization using distributed bag-of-words model. In *2018 IEEE/WIC/ACM International Conference on Web Intelligence (WI)*, pp. 672–675. IEEE.
- Marcos, Diego, V. (2016). Learning rotation invariant convolutional filters for texture classification. In *2016 23rd International Conference on Pattern Recognition*, pp. 2012–2017. IEEE.
- Meister, C., Vieira, T., & Cotterell, R. (2020). If beam search is the answer, what was the question?. *Preprint arXiv:2010.02650, ArXiv*.
- Mena, G., Belanger, D., Linderman, S., & Snoek, J. (2018). Learning latent permutations with gumbel-sinkhorn networks. In *International Conference on Learning Representations 2018*.
- Meng, C., Yang, J., Ribeiro, B., & Neville, J. (2019). HATS: A hierarchical sequence-attention framework for inductive set-of-sets embeddings. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2019*, pp. 783–792.
- Merity, S., Xiong, C., Bradbury, J., & Socher, R. (2016). Pointer sentinel mixture models. In *5th International Conference on Learning Representations, ICLR 2017 - Conference Track Proceedings*.
- Mitra, B., Craswell, N., et al. (2018). An introduction to neural information retrieval. *Foundations and Trends® in Information Retrieval*, 13(1), 1–126.
- Moore, J., & Neville, J. (2017). Deep collective inference. In *AAAI*.
- Murphy, R. L., Srinivasan, B., Ribeiro, B., & Rao, V. (2019). Janossy pooling: Learning deep permutation-invariant functions for variable-size inputs. In *7th International Conference on Learning Representations, ICLR 2019*, pp. 1–21.
- Nazari, M., Oroojlooy, A., Takáč, M., & Snyder, L. V. (2018). Reinforcement learning for solving the vehicle routing problem. *Advances in Neural Information Processing Systems, 2018-December*, 9839–9849.
- Niepert, M., Ahmed, M., & Kutzkov, K. (2016). Learning convolutional neural networks for graphs. In *International conference on machine learning*, pp. 2014–2023.

- Nishida, N., & Nakayama, H. (2017). Word ordering as unsupervised learning towards syntactically plausible word representations. In *Proceedings of the 8th International Joint Conference on Natural Language Processing (IJCNLP 2017)*, pp. 70–79.
- Noreen, N., Palaniappan, S., Qayyum, A., Ahmad, I., Imran, M., & Shoaib, M. (2020). A deep learning model based on concatenation approach for the diagnosis of brain tumor. *IEEE Access*, 8, 55135–55144.
- Ntampaka, M., Trac, H., Sutherland, D., Fromenteau, S., Poczoz, B., & Schneider, J. (2016). Dynamical mass measurements of contaminated galaxy clusters using machine learning. *The Astrophysical Journal*, 831.
- Oladosu, A., Xu, T., Ekfeldt, P., Kelly, B. A., Cranmer, M., Ho, S., Price-Whelan, A. M., & Contardo, G. (2020). Meta-learning one-class classification with deepsets: Application in the milky way. *Preprint arXiv:2007.04459*, *ArXiv*.
- Oliveira, H., Silva, C., Machado, G. L., Nogueira, K., & Santos, J. A. d. (2020). Fully convolutional open set segmentation. *Preprint arXiv:2006.14673*, *ArXiv*.
- Pang, L., Xu, J., Ai, Q., Lan, Y., Cheng, X., & Wen, J. (2020). Setrank: Learning a permutation-invariant ranking model for information retrieval. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 499–508.
- Pfannschmidt, K., Gupta, P., & Hüllermeier, E. (2018). Deep architectures for learning context-dependent ranking functions. *ArXiv, abs/1803.05796*.
- Pihera, J., & Musliu, N. (2014). Application of machine learning to algorithm selection for tsp. In *2014 IEEE 26th International Conference on Tools with Artificial Intelligence*, pp. 47–54. IEEE.
- Puduppully, R., Dong, L., & Lapata, M. (2019). Data-to-Text Generation with Content Selection and Planning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33, 6908–6915.
- Qi, C. R., Su, H., Mo, K., & Guibas, L. J. (2017). PointNet: Deep learning on point sets for 3D classification and segmentation. In *Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017*, pp. 77–85.
- Rang, B., & Thomas, W. (1981). Zermelo’s discovery of the “russell paradox”. *Historia Mathematica*, 8(1), 15–22.
- Rocktäschel, T., Grefenstette, E., Hermann, K., Kociský, T., & Blunsom, P. (2016). Reasoning about entailment with neural attention. *CoRR, abs/1509.06664*.
- Rolínek, M., Musil, V., Paulus, A., Vlastelica, M., Michaelis, C., & Martius, G. (2020). Optimizing rank-based metrics with blackbox differentiation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7620–7630.
- Saito, Y., Nakamura, T., Hachiya, H., & Fukumizu, K. (2019). Deep set-to-set matching and learning. *ArXiv, abs/1910.09972*.
- Sarwar, B., Karypis, G., Konstan, J., & Riedl, J. (2001). Item-based collaborative filtering recommendation algorithms. In *Proceedings of the 10th International Conference on World Wide Web, WWW '01*, p. 285–295, New York, NY, USA. Association for Computing Machinery.
- Schafer (2007). Recurrent neural networks are universal approximators. *International Journal of Neural Systems, Vol. 17, No. 4 (2007) 253–263*, World Scientific Publishing Company Int., 17(4), 253–263.
- See, A., Liu, P., & Manning, C. (2017). Get to the point: Summarization with pointer-generator networks. In *Association for Computational Linguistics*.
- Serviansky, H., Segol, N., Shlomi, J., Cranmer, K., Gross, E., Maron, H., & Lipman, Y. (2020). Set2graph: Learning graphs from sets. *Preprint arXiv:2002.08772*, *ArXiv*.

- Skianis, K., Nikolentzos, G., Limmios, S., & Vazirgiannis, M. (2020). Rep the Set: Neural Networks for Learning Set Representations. *AISTATS*, 1(23rd).
- Smith, K. A. (1999). Neural networks for combinatorial optimization: A review of more than a decade of research. *INFORMS Journal on Computing*, 11(1), 15–34.
- Spearman, C. (1904). The proof and measurement of association between two things.. In *American Journal of Psychology*.
- Stuart, J. L., & Weaver, J. R. (1991). Matrices that commute with a permutation matrix. *Linear Algebra and its Applications*, 150, 255 – 265.
- Sun, Z., Tang, J., Du, P., Deng, Z. H., & Nie, J. Y. (2019). DivGraphPointer: A graph pointer network for extracting diverse keyphrases. In *SIGIR 2019 - Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 755–764.
- Sunehag, P., Lever, G., Gruslys, A., Czarnecki, W. M., Zambaldi, V., Jaderberg, M., Lanctot, M., Sonnerat, N., Leibo, J. Z., Tuyls, K., & Graepel, T. (2018). Value-decomposition networks for cooperative multi-agent learning based on team reward. *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS*, 3, 2085–2087.
- Sutskever, I., Vinyals, O., & Le, Q. V. (2014). Sequence to sequence learning with neural networks. *Advances in Neural Information Processing Systems*, 4(January), 3104–3112.
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., & Wojna, Z. (2016). Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2818–2826.
- Takeuti, G., & Zaring, W. M. (2013). *Axiomatic set theory*, Vol. 8. Springer Science & Business Media.
- Taylor, L., & Nitschke, G. (2018). Improving deep learning with generic data augmentation. In *2018 IEEE Symposium Series on Computational Intelligence (SSCI)*, pp. 1542–1547.
- Vartak, M., & Thiagarajan, A. (2017). A Meta-Learning Perspective on Cold-Start Recommendations for Items Manasi. *31st Conference on Neural Information Processing Systems (NIPS 2017)*, 29(3), 171–180.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems, 2017-Decem(Nips)*, 5999–6009.
- Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., & Bengio, Y. (2018). Graph Attention Networks. In *International Conference on Learning Representations (2018)*. Accepted as poster.
- Vinyals, O., Babuschkin, I., Czarnecki, W. M., & Mathieu (2019). Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, 575(7782), 350–354.
- Vinyals, O., Bengio, S., & Kudlur, M. (2016). Order matters: Sequence to sequence for sets. In *4th International Conference on Learning Representations, ICLR 2016 - Conference Track Proceedings*, pp. 1–11.
- Vinyals, O., Fortunato, M., & Jaitly, N. (2015). Pointer networks. In *Advances in Neural Information Processing Systems*, Vol. 2015-January, pp. 2692–2700.
- Volkovs, M. N., & Zemel, R. S. (2009). Boltzrank: learning to maximize expected ranking gain. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pp. 1089–1096.
- Wagstaff, E., Fuchs, F. B., Engelcke, M., Posner, I., & Osborne, M. (2019). On the limitations of representing functions on sets. *36th International Conference on Machine Learning, ICML 2019, 2019-June*, 11285–11298.

- Wang, J., Song, Y., Leung, T., Rosenberg, C., Wang, J., Philbin, J., Chen, B., & Wu, Y. (2014). Learning fine-grained image similarity with deep ranking. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 1386–1393.
- Wang, T., & Wan, X. (2019). Hierarchical Attention Networks for Sentence Ordering. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33, 7184–7191.
- Weston, J., Chopra, S., & Bordes, A. (2015). Memory Networks. In *International Conference on Learning Representations, ICLR, 2015*, pp. 1–15.
- Williams, R. J. (1992). Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4), 229–256.
- Wiseman, S., & Rush, A. M. (2016). Sequence-to-sequence learning as beam-search optimization. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 1296–1306, Austin, Texas. Association for Computational Linguistics.
- Wolf, T., Sanh, V., Chaumond, J., & Delangue, C. (2019). TransferTransfo: A Transfer Learning Approach for Neural Network Based Conversational Agents. *Association for the Advancement of Artificial Intelligence*, 1(ii).
- Wu, Y., Riedel, S., Minervini, P., & Stenetorp, P. (2020). Don’t read too much into it: Adaptive computation for open-domain question answering. *Preprint arXiv:2011.05435, ArXiv*.
- Wu, Z., Song, S., Khosla, A., Yu, F., Zhang, L., Tang, X., & Xiao, J. (2015). 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1912–1920.
- Xia, F., Liu, T.-Y., Wang, J., Zhang, W., & Li, H. (2008). Listwise approach to learning to rank: theory and algorithm. In *Proceedings of the 25th international conference on Machine learning*, pp. 1192–1199.
- Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A., Salakhudinov, R., Zemel, R., & Bengio, Y. (2015). Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*, pp. 2048–2057.
- Xu, K., Jegelka, S., Hu, W., & Leskovec, J. (2019). How powerful are graph neural networks?. In *7th International Conference on Learning Representations, ICLR 2019*, pp. 1–17.
- Xu, Z., Zhu, L., & Yang, Y. (2017). Few-shot object recognition from machine-labeled web images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1164–1172.
- Yang, B., Wang, S., Markham, A., & Trigoni, N. (2020). Robust attentional aggregation of deep feature sets for multi-view 3d reconstruction. *International Journal of Computer Vision*, 128(1), 53–73.
- Ying, Z., You, J., Morris, C., Ren, X., Hamilton, W., & Leskovec, J. (2018). Hierarchical graph representation learning with differentiable pooling. In *Advances in neural information processing systems*, pp. 4800–4810.
- Yu, T., Meng, J., & Yuan, J. (2018). Multi-view harmonized bilinear network for 3d object recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 186–194.
- Zaheer, M., Kottur, S., Ravanbakhsh, S., Póczos, B., Salakhutdinov, R., & Smola, A. J. (2017). Deep sets. *Advances in Neural Information Processing Systems, 2017-December(ii)*, 3392–3402.
- Zhang, M., Cui, Z., Neumann, M., & Chen, Y. (2018). An end-to-end deep learning architecture for graph classification. In *AAAI*.
- Zhang, Y., Hare, J., & Prügel-Bennett, A. (2019). Deep Set Prediction Networks. In *NeurIPS 2019*.

- Zhang, Y., Hare, J., & Prügel-Bennett, A. (2020). FSPool: Learning Set Representations with Featurewise Sort Pooling. In *ICLR 2020*.
- Zhang, Y., Prügel-Bennett, A., & Hare, J. (2019). Learning representations of sets through optimized permutations. In *7th International Conference on Learning Representations, ICLR 2019*.
- Zhao, Z.-Q., tao Xu, S., Liu, D., Tian, W., & Jiang, Z.-D. (2019a). A review of image set classification. *Neurocomputing*, *335*, 251–260.
- Zhao, Z. Q., Zheng, P., Xu, S. T., & Wu, X. (2019b). Object Detection with Deep Learning: A Review..
- Zheng, L., Shen, L., Tian, L., Wang, S., Wang, J., & Tian, Q. (2015). Scalable person re-identification : A benchmark university of texas at san antonio. In *ICCV 2015*, pp. 1116–1124.
- Zhong, V., Xiong, C., & Socher, R. (2018). Seq2sql: Generating structured queries from natural language using reinforcement learning. In *6th International Conference on Learning Representations, ICLR 2018 - Conference Track Proceedings*.