Dear Author,

Here are the proofs of your article.

- You can submit your corrections online, via e-mail or by fax.
- For **online** submission please insert your corrections in the online correction form. Always indicate the line number to which the correction refers.
- You can also insert your corrections in the proof PDF and email the annotated PDF.
- For fax submission, please ensure that your corrections are clearly legible. Use a fine black pen and write the correction in the margin, not too close to the edge of the page.
- Remember to note the **journal title**, **article number**, and **your name** when sending your response via email or fax.
- **Check** the metadata sheet to make sure that the header information, especially author names and the corresponding affiliations are correctly shown.
- Check the questions that may have arisen during copy editing and insert your answers/ corrections.
- **Check** that the text is complete and that all figures, tables and their legends are included. Also check the accuracy of special characters, equations, and electronic supplementary material if applicable. If necessary refer to the *Edited manuscript*.
- The publication of inaccurate data such as dosages and units can have serious consequences. Please take particular care that all such details are correct.
- Please **do not** make changes that involve only matters of style. We have generally introduced forms that follow the journal's style. Substantial changes in content, e.g., new results, corrected values, title and authorship are not allowed without the approval of the responsible editor. In such a case, please contact the Editorial Office and return his/her consent together with the proof.
- If we do not receive your corrections within 48 hours, we will send you a reminder.
- Your article will be published **Online First** approximately one week after receipt of your corrected proofs. This is the **official first publication** citable with the DOI. **Further changes are, therefore, not possible.**
- The **printed version** will follow in a forthcoming issue.

Please note

After online publication, subscribers (personal/institutional) to this journal will have access to the complete article via the DOI using the URL: http://dx.doi.org/[DOI].

If you would like to know when your article has been published online, take advantage of our free alert service. For registration and further information go to: <u>http://www.link.springer.com</u>.

Due to the electronic nature of the procedure, the manuscript and the original figures will only be returned to you on special request. When you return your corrections, please inform us if you would like to have these documents returned.

Metadata of the article that will be visualized in OnlineFirst

ArticleTitle	Learning the travelling	g salesperson problem requires rethinking generalization
Article Sub-Title		
Article CopyRight	The Author(s) (This will be the copy	right line in the final PDF)
Journal Name	Constraints	
Corresponding Author	FamilyName Particle	Joshi
	Given Name Suffix Division	Chaitanya K.
	Organization Address Phone	University of Cambridge Cambridge, UK
	Fax Email URL	ckj24@cl.cam.ac.uk
	ORCID	http://orcid.org/0000-0003-4722-1815
Author	FamilyName Particle	Cappart
	Given Name Suffix Division	Quentin
	Organization Address Phone Fax	Ecole Polytechnique de Montréal Montreal, Canada
	Email URL ORCID	quentin.cappart@polymtl.ca
Author	FamilyName Particle	Rousseau
	Given Name Suffix Division	Louis-Martin
	Organization Address Phone Fax	Ecole Polytechnique de Montréal Montreal, Canada
	Email URL ORCID	louis-martin.rousseau@polymtl.ca
Author	FamilyName Particle	Laurent
	Given Name Suffix Division	Thomas
	Organization	Loyola Marymount University

	Address Phone Fax	Los Angeles, USA	
	Email URL ORCID	tlaurent@lmu.edu	
Schedule	Received Revised Accepted	16 Mar 2022	
Abstract	End-to-end training of neural network solvers for graph combinatorial optimization problems such as the Travelling Salesperson Problem (TSP) have seen a surge of interest recently, but remain intractable and inefficient beyond graphs with few hundreds of nodes. While state-of-the-art learning-driven approaches for TSP perform closely to classical solvers when trained on trivially small sizes, they are unable to generalize the learnt policy to larger instances at practical scales. This work presents an end-to-end <i>neural combinatorial optimization</i> pipeline that unifies several recent papers in order to identify the inductive biases, model architectures and learning algorithms that promote generalization to instances larger than those seen in training. Our controlled experiments provide the first principled investigation into such <i>zero-shot</i> generalization, revealing that extrapolating beyond training data requires rethinking the neural combinatorial optimization pipeline, from network layers and learning paradigms to evaluation protocols. Additionally, we analyze recent advances in deep learning for routing problems through the lens of our pipeline and provide new directions to stimulate future research.		
Keywords (separated by '-')	Combinatorial optimizatio	on - Travelling salesperson problem - Graph neural networks - Deep learning	
Footnote Information	Code and datasets: github	.com/chaitjo/learning-tsp	

Journal	:	SmallCondensed	10601	1
---------	---	----------------	-------	---

Article No : 9327



Learning the travelling salesperson problem requiresrethinking generalization

³ Chaitanya K. Joshi¹ · Quentin Cappart² · Louis-Martin Rousseau² · Thomas Laurent³

Accepted: 16 March 2022

⁵ © The Author(s) 2022

6 Abstract

7 End-to-end training of neural network solvers for graph combinatorial optimization 8 problems such as the Travelling Salesperson Problem (TSP) have seen a surge of inter-9 est recently, but remain intractable and inefficient beyond graphs with few hundreds of 10 nodes. While state-of-the-art learning-driven approaches for TSP perform closely to clas-11 sical solvers when trained on trivially small sizes, they are unable to generalize the learnt 12 policy to larger instances at practical scales. This work presents an end-to-end neural com-13 binatorial optimization pipeline that unifies several recent papers in order to identify the 14 inductive biases, model architectures and learning algorithms that promote generalization 15 to instances larger than those seen in training. Our controlled experiments provide the first 16 principled investigation into such zero-shot generalization, revealing that extrapolating 17 beyond training data requires rethinking the neural combinatorial optimization pipeline, 18 from network layers and learning paradigms to evaluation protocols. Additionally, we ana-19 lyze recent advances in deep learning for routing problems through the lens of our pipeline 20 and provide new directions to stimulate future research.

²¹ Keywords Combinatorial optimization · Travelling salesperson problem · Graph neural

- 22 networks · Deep learning
- A1 Code and datasets: github.com/chaitjo/learning-tsp

A2	\square	Chaitanya K. Joshi	i
~~			

A3 ckj24@cl.cam.ac.uk

- A4 Quentin Cappart A5 quentin.cappart@polymtl.ca
- A6 Louis-Martin Rousseau A7 louis-martin.rousseau@polymtl.ca

A8 Thomas Laurent

- A9 tlaurent@lmu.edu
- A10¹ University of Cambridge, Cambridge, UK
- A11² Ecole Polytechnique de Montréal, Montreal, Canada
- A12 ³ Loyola Marymount University, Los Angeles, USA

Iournal · SmallCondensed 10601

Constraints

1 Introduction 23

NP-hard combinatorial optimization problems are the family of integer constrained optimiza-24 tion problems which are intractable to solve optimally at large scales. Robust approximation 25 algorithms to popular problems have immense practical applications and are the backbone of 26 modern industries. Among combinatorial problems, the 2D Euclidean Travelling Salesperson AQ1 27 Problem (TSP) has been the most intensely studied NP-hard graph problem in the Opera-28 29 tions Research (OR) community, with applications in logistics, genetics and scheduling [1]. TSP is intractable to solve optimally above thousands of nodes for modern computers [2]. In 30 practice, the Concorde TSP solver [3] uses linear programming with carefully handcrafted 31 heuristics to find solutions up to tens of thousands of nodes, but with prohibitive execution 32 times.¹ Besides, the development of problem-specific OR solvers such as Concorde for novel 33 34 or under-studied problems arising in scientific discovery [4] or computer architecture [5] AQ2 requires significant time and specialized knowledge. 35 An alternate approach by the Machine Learning community is to develop generic learning 36 algorithms which can be trained to solve any combinatorial problem directly from problem 37 instances themselves [6-8]. Using classical problems such as TSP, Minimum Vertex Cover and 38

Boolean Satisfiability as benchmarks, recent end-to-end approaches [9–11] leverage advances 39 40 in graph representation learning [12-15] and have shown competitive performance with OR solvers on trivially small problem instances up to few hundreds of nodes. Once trained, approxi-41

mate solvers based on Graph Neural Networks (GNNs) have significantly favorable time com-42

43 plexity than their OR counterparts, making them highly desirable for real-time decision-making

problems such as TSP and the associated class of Vehicle Routing Problems (VRPs). 44

1.1 Motivation 45

Scaling end-to-end approaches to practical and real-world instances is still an open question [8] 46 as the training phase of state-of-the-art models on large graphs is extremely time-consuming. 47 For graphs larger than few hundreds of nodes, the gap between GNN-based solvers and simple 48 non-learnt heuristics is especially evident for routing problems like TSP [16, 17]. 49

As an illustration, Fig. 1 presents the computational challenge of learning TSP on 200-50 node graphs (TSP200) in terms of both sample efficiency and wall clock time. Surpris-51 ingly, it is difficult to outperform a simple insertion heuristic when directly training on 12.8 52 53 Million TSP200 samples for 500 hours on university-scale hardware.

54 We advocate for an alternative to expensive large-scale training: learning efficiently from trivially small TSP and transferring the learnt policy to larger graphs in a zero-shot 55 fashion or via fast finetuning. Thus, identifying promising inductive biases, architectures 56 and learning paradigms that enable such zero-shot generalization to large and more com-57 plex instances is a key concern for training practical solvers for real-world problems. 58

1.2 Contributions 59

Towards end-to-end learning of scale-invariant TSP solvers, we unify several state-of-60 the-art architectures and learning paradigms [16-19] into one experimental pipeline and 61

Article No : 9327

Dispatch : 21-4-2022

¹ The largest TSP solved by Concorde to date has 109,399 nodes with running time of 7.5 months.

Journal : SmallCondensed 10601	Article No : 9327	Pages : 30	MS Code : 9327	Dispatch : 21-4-2022
--------------------------------	-------------------	------------	----------------	----------------------



Fig. 1 Computational challenges of learning large scale TSP. We compare three identical autoregressive GNN-based models trained on 12.8 Million TSP instances via reinforcement learning. We plot average optimality gap to the Concorde solver on 1,280 held-out TSP200 instances vs. number of training samples (left) and wall clock time (right) during the learning process. Training on large TSP200 from scratch is intractable and sample inefficient. Active Search [7], which learns to directly overfit to the 1,280 held-out samples, further demonstrates the computational challenge of memorizing very few TSP200 instances. Comparatively, learning efficiently from trivial TSP20-TSP50 allows models to better generalize to TSP200 in a zero-shot manner, indicating positive knowledge transfer from small to large graphs. Performance can further improve via rapid finetuning on 1.28 Million TSP200 instances or by Active Search. Within our computational budget, a simple non-learnt *furthest insertion* heuristic still outperforms all models. Precise experimental setup is described in Appendix A

62 provide the first principled investigation on zero-shot generalization to large instances. Our

63 findings suggest that learning scale-invariant TSP solvers requires rethinking the status quo

64 of neural combinatorial optimization to explicitly account for generalization:

- The prevalent evaluation paradigm overshadows models' poor generalization capabilities by measuring performance on fixed or trivially small TSP sizes.
- Generalization performance of GNN aggregation functions and normalization schemes
 benefits from explicit redesigns which account for shifting graph distributions, and can

be further boosted by enforcing regularities such as constant graph diameters whendefining problems using graphs.

- Autoregressive decoding enforces a sequential inductive bias which improves generali zation over non-autoregressive models, but is costlier in terms of inference time.
- Models trained with expert supervision are more amenable to post-hoc search, while
 reinforcement learning approaches scale better with more computation as they do not
 rely on labelled data.

Our framework and datasets are available online.² Additionally, we use our pipeline to characterize the recent state-of-the-art in deep learning for routing problems and provide new directions for future research.

79 2 Related work

80 **Neural combinatorial optimization** In a recent survey, Bengio et al. [8] identified three 81 broad approaches to leveraging machine learning for combinatorial optimization problems:

82 learning alongside optimization algorithms [20-22], learning to configure optimization

²FL01² https://github.com/chaitjo/learning-tsp

algorithms [23, 24], and end-to-end learning to approximately solve optimization problems, *a.k.a.* neural combinatorial optimization [6, 7].

State-of-the-art end-to-end approaches for TSP use Graph Neural Networks (GNNs) [12–15] and *sequence-to-sequence* learning [25] to construct approximate solutions directly from problem instances. Architectures for TSP can be classified as: (1) autoregressive approaches, which build solutions in a step-by-step fashion [9, 16, 19, 26–28]; and (2) non-autoregressive models, which produce the solution in one shot [17, 18, 29–31]. Models can be trained to imitate optimal solvers via supervised learning or by minimizing the length of TSP tours via reinforcement learning [32].

Other classical problems tackled by similar architectures include Vehicle Routing [33, 34], Maximum Cut [9], Minimum Vertex Cover [11], Boolean Satisfiability [10, 35], and Graph Coloring [36]. Using TSP as an illustration, we present a unified pipeline for charac-

95 terizing neural combinatorial optimization architectures in Section 3.

Notably, TSP has emerged as a challenging testbed for neural combinatorial optimization. Whereas generalization to problem instances larger and more complex than those seen in training has at least partially been demonstrated on non-sequential problems such as SAT, MaxCut, and MVC [10, 11], the same architectures do not show strong generalization for TSP [16, 17].

Combinatorial optimization and GNNs From the perspective of graph representa-101 tion learning, algorithmic and combinatorial problems have recently been used to 102 characterize the expressive power of GNNs [37, 38]. An emerging line of work on 103 learning to execute graph algorithms [39, 40] has lead to the development of prov-104 ably more expressive GNNs [41] and improved understanding of their generalization 105 capability [42, 43]. Towards tackling realistic and large-scale combinatorial prob-106 lems, this paper aims to quantify the limitations of prevalent GNN architectures and 107 learning paradigms via zero-shot generalization to problems larger than those seen 108 during training. 109

Novel applications Advances on classical combinatorial problems have shown promising results in downstream applications to novel or under-studied optimization problems in the physical sciences [4, 44] and computer architecture [5, 45, 46], where the development of exact solvers is expensive and intractable. For example, autoregressive architectures provide a strong inductive bias for device placement optimization problems [47, 48], while non-autoregressive models [49] are competitive with autoregressive approaches [50, 51] for molecule generation tasks.

117 3 Neural combinatorial optimization pipeline

118 NP-hard problems can be formulated as sequential decision making tasks on graphs 119 due to their highly structured nature. Towards a controlled study of neural combinato-120 rial optimization for TSP, we unify recent ideas [16–19] via a five stage end-to-end 121 pipeline illustrated in Fig. 2. Our discussion focuses on TSP, but the pipeline presented 122 is generic and can be extended to characterize modern architectures for several NP-123 hard graph problems.



(a) Neural combinatorial optimization pipeline in stages.



(b) **Problem Definition:** TSP is formulated via a fully-connected graph of cities/nodes. The graph can be sparsified via heuristics such as k-nearest neighbors.



(c) **Graph Embedding:** Embeddings for each graph node are obtained using a Graph Neural Network encoder, which builds local structural features.



(d) **Solution Decoding & Search:** Probabilities are assigned to each node for belonging to the solution set, either independent of one-another (*i.e.* Non-autoregressive decoding) or conditionally through graph traversal (*i.e.* Autoregressive decoding). The predicted probabilities are converted into discrete decisions through classical graph search techniques such as greedy search or beam search.

Fig. 2 End-to-end neural combinatorial optimization pipeline: The entire model in trained end-to-end via imitating an optimal solver (i.e. supervised learning) or through minimizing a cost function (i.e. reinforcement learning)

Journal : SmallCondensed 10601	Article No : 9327	Pages : 30	MS Code : 9327	Dispatch : 21-4-2022
--------------------------------	-------------------	------------	----------------	----------------------

124 3.1 Problem definition

The 2D Euclidean TSP is defined as follows: "Given a set of cities and the distances between each pair of cities, what is the shortest possible route that visits each city and returns to the origin city?" Formally, given a fully-connected input graph of *n* cities (nodes) in the two dimensional unit square $S = \{x_i\}_{i=1}^n$ where each $x_i \in [0, 1]^2$, we aim to find a permutation of the nodes π , termed a tour, that visits each node once and has the minimum total length, defined as:

131

$$L(\pi|s) = \|x_{\pi_n} - x_{\pi_1}\|_2 + \sum_{i=1}^{n-1} \|x_{\pi_i} - x_{\pi_{i+1}}\|_2,$$
(1)

132

133 where $\|\cdot\|_2$ denotes the ℓ_2 norm.

Graph sparsification Classically, TSP is defined on fully-connected graphs, see Fig. 2(b). Graph sparsification heuristics based on *k*-nearest neighbors aim to reduce TSP graphs, enabling models to scale up to large instances where pairwise computation for all nodes is intractable [9] or learn faster by reducing the search space [17]. Notably, problem-specific graph reduction techniques have proven effective for out-of-distribution generalization to larger graphs for other NP-hard problems such as MVC and SAT [11].

Fixed size vs. variable size graphs Most work on learning for TSP has focused on training with a fixed graph size [16, 17], likely due to ease of implementation. Learning from multiple graph sizes naturally enables better generalization within training size ranges, but its impact on generalization to larger TSP instances remains to be analyzed.

144 3.2 Graph embedding

A Graph Neural Network (GNN) encoder computes *d*-dimensional representations for each node in the input TSP graph, see Fig. 2(c). At each layer, nodes gather features from their neighbors to represent local graph structure via recursive message passing [13]. Stacking *L* layers allows the network to build representations from the *L*-hop neighborhood of each node. Let h_i^{ℓ} and e_{ij}^{ℓ} denote respectively the node and edge feature at layer ℓ associated with node *i* and edge *ij*. We define the feature at the next layer via an *anisotropic* message pass-AQ3 ing scheme using an edge gating mechanism [52]:

$$h_i^{\ell+1} = h_i^{\ell} + \operatorname{ReLU}\left(\operatorname{NORM}\left(U^{\ell}h_i^{\ell} + \operatorname{AGGR}_{j \in \mathcal{N}_i}\left(\sigma(e_{ij}^{\ell}) \odot V^{\ell}h_j^{\ell}\right)\right)\right),$$
(2)

$$e_{ij}^{\ell+1} = e_{ij}^{\ell} + \operatorname{ReLU}\left(\operatorname{NORM}\left(A^{\ell}e_{ij}^{\ell} + B^{\ell}h_{i}^{\ell} + C^{\ell}h_{j}^{\ell}\right)\right),\tag{3}$$

155

where $U^{\ell}, V^{\ell}, A^{\ell}, B^{\ell}, C^{\ell} \in \mathbb{R}^{d \times d}$ are learnable parameters, NORM denotes the normalization layer (BatchNorm [53], LayerNorm [54]), AGGR represents the neighborhood aggregation function (SUM, MEAN or MAX), σ is the sigmoid function, and \odot is the Hadamard product. As inputs $h_i^{\ell=0}$ and $e_{ij}^{\ell=0}$, we use *d*-dimensional linear projections of the node coordinate x_i and the euclidean distance $||x_i - x_j||_2$, respectively.

Journal : SmallCondensed 10601	Article No : 9327	Pages : 30	MS Code : 9327	Dispatch : 21-4-2022	1
--------------------------------	-------------------	------------	----------------	----------------------	---

161 Anisotropic aggregation We make the aggregation function anisotropic or directional 162 via a dense attention mechanism which scales the neighborhood features h_j , $\forall j \in \mathcal{N}_i$, using 163 edge gates $\sigma(e_{ij})$. Anisotropic and attention-based GNNs such as Graph Attention Networks 164 [14], Transformers [55, 56], and Gated Graph ConvNets [52] have been shown to outper-165 form isotropic Graph ConvNets [12] across several challenging domains [57], including 166 TSP [16, 17].

167 3.3 Solution decoding

168 **Non-autoregressive decoding (NAR)** Consider TSP as a link prediction task: each edge 169 may belong/not belong to the optimal TSP solution independent of one another [18]. We 170 define the edge predictor as a two layer MLP on the node embeddings produced by the 171 final GNN encoder layer L, following Joshi et al. [17], see Fig. 2(d). For adjacent nodes *i* 172 and *j*, we compute the unnormalized edge logits:

173

$$\hat{p}_{ij} = W_2 \Big(\text{ReLU} \Big(W_1 \Big(\Big[h_G, h_i^L, h_j^L \Big] \Big) \Big) \Big), \text{ where } h_G = \frac{1}{n} \sum_{i=0}^n h_i^L, \tag{4}$$

174

175 $W_1 \in \mathbb{R}^{3d \times d}, W_2 \in \mathbb{R}^{d \times 2}$, and $[\cdot, \cdot, \cdot]$ is the concatenation operator. The logits \hat{p}_{ij} are con-176 verted to probabilities over each edge p_{ij} via a softmax.

Autoregressive decoding (AR) Although NAR decoders are fast as they produce predic-177 tions in one shot, they ignore the sequential ordering of TSP tours. Autoregressive decod-178 ers, based on attention [16, 19] or recurrent neural networks [6, 26], explicitly model this 179 sequential inductive bias through step-by-step graph traversal. We follow the attention 180 decoder from Kool et al. [16], which starts from a random node and outputs a probability 181 distribution over its neighbors at each step. Greedy search is used to perform the traversal 182 over *n* time steps and masking enforces constraints such as not visiting previously visited 183 nodes. 184

At time step *t* at node *i*, the decoder builds a context \hat{h}_i^C for the partial tour $\pi'_{t'}$, generated at time t' < t, by packing together the graph embedding h_G and the embeddings of the first and last node in the partial tour: $\hat{h}_i^C = W_C \left[h_G, h_{\pi'_{t-1}}^L, h_{\pi'_1}^L \right]$, where $W_C \in \mathbb{R}^{3d \times d}$ and learned placeholders are used for $h_{\pi'_{t-1}}^L$ and $h_{\pi'_1}^L$ at t = 1. The context \hat{h}_i^C is then refined via a standard Multi-Head Attention (MHA) operation [55] over the node embeddings:

$$h_i^C = \text{MHA}(Q = \hat{h}_i^C, K = \{h_1^L, \dots, h_n^L\}, V = \{h_1^L, \dots, h_n^L\}),$$
(5)

where Q, K, V are inputs to the *M*-headed MHA (M = 8). The unnormalized logits for each edge e_{ij} are computed via a final attention mechanism between the context h_i^C and the mbedding h_j :

$$\hat{p}_{ij} = \begin{cases} C \cdot \tanh\left(\frac{\left(W_{\mathcal{Q}}h_i^c\right)^T \cdot \left(W_{\mathcal{K}}h_j^L\right)}{\sqrt{d}}\right) & \text{if } j \neq \pi_{t'} \quad \forall t' < t\\ -\infty & \text{otherwise.} \end{cases}$$
(6)

196

197 The tanh is used to maintain the value of the logits within [-C,C] (C = 10) [7]. The logits 198 \hat{p}_{ii} at the current node *i* are converted to probabilities p_{ii} via a softmax over all edges.

Journal : SmallCondensed 10601 Article No : 9327 Pages : 30 MS Code : 9327 Dispatch : 21-4-2022

Inductive biases NAR approaches, which make predictions over edges independently of one-another, have shown strong out-of-distribution generalization for non-sequential problems such as SAT and MVC [11]. On the other hand, AR decoders come with the sequential/tour constraint built-in and are the default choice for routing problems [16]. Although both approaches have shown close to optimal performance on fixed and small TSP sizes under different experimental settings, it is important to fairly compare which inductive biases are most useful for generalization.

206 3.4 Solution search

Greedy search For AR decoding, the predicted probabilities at node *i* are used to select the edge to travel along at the current step via sampling from the probability distribution p_i or greedily selecting the most probable edge p_{ij} , i.e. greedy search. Since NAR decoders directly output probabilities over all edges independent of one-another, we can obtain valid TSP tours using greedy search to traverse the graph starting from a random node and masking previously visited nodes. Thus, the probability of a partial tour π' can be formulated as $p(\pi') = \prod_{i' \sim i' \in \pi'} p_{i'i'}$, where each node j' follows node i'.

Beam search and sampling During inference, we can increase the capacity of greedy 214 search via limited width breadth-first beam search, which maintains the b most probable 215 tours during decoding. Similarly, we can sample b solutions from the learnt policy and 216 select the shortest tour among them. Naturally, searching longer, with more sophisticated 217 techniques, or sampling more solutions allows trading off run time for solution quality. 218 However, it has been noted that using large b for search/sampling or local search during 219 inference may overshadow an architecture's inability to generalize [58]. To better under-220 stand generalization, we focus on using greedy search and beam search/sampling with 221 small b = 128. 222

223 3.5 Policy learning

Supervised learning Models can be trained end-to-end via imitating an optimal solver at each step (i.e. supervised learning). For models with NAR decoders, the edge predictions are linked to the ground-truth TSP tour by minimizing the binary cross-entropy loss for each edge [17]. For AR architectures, at each step, we minimize the cross-entropy loss between the predicted probability distribution over all edges leaving the current node and the next node from the groundtruth tour, following Vinyals et al. [6]. We use teacher-forcing to stabilize training [59].

Reinforcement learning Reinforcement learning is a elegant alternative in the absence of 231 groundtruth solutions, as is often the case for understudied combinatorial problems. Mod-232 els can be trained by minimizing problem-specific cost functions (the tour length in the 233 case of TSP) via policy gradient algorithms [7, 16] or Q-Learning [9]. We focus on policy 234 gradient methods due to their simplicity, and define the loss for an instance s parameterized 235 236 by the model θ as $\mathcal{L}(\theta|s) = \mathbb{E}_{p_a(\pi|s)}[L(\pi)]$, the expectation of the tour length $L(\pi)$, where $p_{\theta}(\pi | s)$ is the probability distribution from which we sample to obtain the tour $\pi | s$. We use 237 the REINFORCE gradient estimator [60] to minimize \mathcal{L} : 238

Journal : SmallCondensed 10601	Article No : 9327	Pages : 30	MS Code : 9327	Dispatch : 21-4-2022
--------------------------------	-------------------	------------	----------------	----------------------

239 240

$$\nabla \mathcal{L}(\theta|s) = \mathbb{E}_{p_{\theta}(\pi|s)} \left[(L(\pi) - b(s)) \nabla \log p_{\theta}(\pi|s) \right], \tag{7}$$

where the baseline b(s) reduces gradient variance. Our experiments compare standard critic network baselines [7, 19] and the greedy rollout baseline proposed by Kool et al. [16].

243 4 Experimental setup

We design controlled experiments to probe the unified pipeline described in Section 3 in 244 order to identify inductive biases, architectures and learning paradigms that promote zero-245 246 shot generalization. We focus on learning efficiently from small problem instances (TSP20-50) and measure generalization to a wider range of sizes, including large instances which 247 are intractable to learn from (e.g. TSP200). Each experiment starts with a 'base' model 248 configuration and ablates the impact of a specific component of the five-stage pipeline. We 249 aim to fairly compare state-of-the-art ideas in terms of model capacity and training data, 250 and expect models with good inductive biases for TSP to: (1) learn trivially small TSPs 251 without hundreds of millions of training samples and model parameters; and (2) generalize 252 reasonably well across smaller and larger instances than those seen in training. 253

To quantify 'good' generalization, we additionally evaluate our models against a simple, non-learnt *furthest insertion* heuristic baseline, which constructively builds a partial tour π' by inserting node *i* between tour nodes $j_1, j_2 \in \pi'$ such that the distance from node *i* to its nearest node j_1 is maximized. Kool et al. [16] provide a detailed description of insertion heuristic baselines.

Training datasets We perform ablation studies of each component of the pipeline by 259 training on variable TSP20-50 graphs for rapid experimentation. We also compare to learn-260 ing from fixed graph sizes up to TSP100. Each TSP instance consist of n nodes sampled 261 uniformly in the unit square $S = \{x_i\}_{i=1}^n$ and $x_i \in [0, 1]^2$. In the supervised learning para-262 digm, we generate a training set of 1,280,000 TSP samples and groundtruth tours using 263 the optimal Concorde solver as an oracle. Models are trained using the Adam optimizer 264 for 10 epochs with a batch size of 128 and a fixed learning rate 1e - 4. For reinforcement 265 learning, models are trained for 100 epochs on 128,000 TSP samples which are randomly 266 generated for each epoch (without optimal solutions) with the same batch size and learn-267 268 ing rate. Thus, both learning paradigms see 12,800,000 TSP samples in total. Considering that TSP20-50 are trivial in terms of complexity as they can be solved by simpler non-269 learnt heuristics, training good solvers at this scale should ideally not require billions of 270 instances. 271

Model hyperparameters For models with AR decoders, we use 3 GNN encoder layers 272 followed by the attention decoder head, setting hidden dimension d = 128. For NAR mod-273 els, we use the same hidden dimension and opt for 4 GNN encoder layers followed by the 274 edge predictor. This results in approximately 350,000 trainable parameters for each model, 275 irrespective of decoder type. Unless specified, most experiments use our best model con-276 figuration: AR decoding scheme and Graph ConvNet encoder with MAX aggregation and 277 BatchNorm (with batch statistics). All models are trained via supervised learning except 278 when comparing learning paradigms. 279

Journal : SmallCondensed 10601	Article No : 9327	Pages : 30	MS Code : 9327	Dispatch : 21-4-2022

Evaluation We compare models on a held-out test set of 25,600 TSPs, consisting of 1,280 280 samples each of TSP10, TSP20, ..., TSP200. Our evaluation metric is the optimality gap 281 w.r.t. the Concorde solver, i.e. the average percentage ratio of predicted tour lengths rela-282 tive to optimal tour lengths. To compare design choices among identical models, we plot 283 line graphs of the optimality gap as TSP size increases (along with a 99%-ile confidence 284 interval) using beam search with a width of 128. Compared to previous work which evalu-285 ated on fixed problem sizes, our protocol identifies not only those models that perform well 286 on training sizes, but also those that generalize better than non-learnt heuristics for large 287 instances which are intractable to train on. 288

289 5 Results

290 5.1 Does learning from variable sizes help generalization?

We train five identical models on fully connected graphs of instances from TSP20, TSP50, 291 TSP100, TSP200 and variable TSP20-50. The line plots of optimality gap across TSP sizes 292 in Fig. 3 indicates that learning from variable TSP sizes helps models retain performance 293 294 across the range of graph sizes seen during training (TSP20-50). Variable graph training compared to training solely on the maximum sized instances (TSP50) leads to marginal 295 gains on small instances but, somewhat counter-intuitively, does not enable better generali-296 zation to larger problems. Learning from small TSP20 is unable to generalize to large sizes 297 while TSP100 models generalize poorly to trivially easy sizes, suggesting that the preva-298 lent protocol of evaluation on training sizes [16, 17] overshadows brittle out-of-distribution 299 performance. 300

Training on TSP200 graphs is intractable within our computational budget, see Fig. 1. TSP100 is the only model which generalizes better to large TSP200 than the non-learnt baseline. However, training on TSP100 can also be prohibitively expensive: one epoch takes approximately 8 hours (TSP100) vs. 2 hours (TSP20-50) (details in Appendix B). For rapid experimentation, we train efficiently on variable TSP20-50 for the rest of our study.



Fig. 3 Learning from various TSP sizes. The prevalent protocol of evaluation on training sizes overshadows brittle out-of-distribution performance to larger and smaller graphs

Journal	:	SmallCondensed	10601	l
---------	---	----------------	-------	---

Article No : 9327 Pag

Constraints

306 5.2 What is the best graph sparsification heuristic?

Figure 4 compares full graph training to the following heuristics: (1) **Fixed node degree** across graph sizes, via connecting each node in TSP*n* to its *k*-nearest neighbors, enabling GNN encoder layers/aggregators to specialize to constant degree *k*; and (2) **Fixed graph diameter** across graph sizes, via connecting each node in TSP*n* to its $n \times k\%$ -nearest neighbors, ensuring that the same number of message passing steps are required to diffuse information across both small and large graphs.

Although both sparsification techniques lead to faster convergence on training instance sizes (not shown), we find that only approach (2) leads to better generalization on larger problems than using full graphs. Consequently, all further experiments use approach (2) to operate on sparse 20%-nearest neighbors graphs. Our results also suggest that developing more principled problem definition and graph reduction techniques beyond simple k-nearest neighbors for augmenting learning-based approaches may be a promising direction.

5.3 What is the relationship between GNN aggregation functionsand normalization layers?

In Fig. 5, we compare identical models with anisotropic SUM, MEAN and MAX aggregation functions. As baselines, we consider the Transformer encoder on full graphs [16, 19] as well as a structure-agnostic MLP on each node, which can be instantiated by not using any aggregation function in (2), i.e. $h_i^{\ell+1} = h_i^{\ell} + \text{ReLU}(\text{NORM}(U^{\ell}h_i^{\ell}))$.

We find that the choice of GNN aggregation function does not have an impact when 325 evaluating models within the training size range TSP20-50. As we tackle larger graphs, 326 GNNs with aggregation functions that are agnostic to node degree (MEAN and MAX) are 327 able to outperform Transformers and MLPs. Importantly, the theoretically more expressive 328 SUM aggregator [61] generalizes worse than structure-agnostic MLPs, as it cannot handle 329 the distribution shift in node degree and neighborhood statistics across graph sizes, lead-330 ing to unstable or exploding node embeddings [39]. We use the MAX aggregator in further 331 experiments, as it generalizes well for both AR and NAR decoders (not shown). 332



Fig. 4 Impact of graph sparsification. Maintaining a constant graph diameter across TSP sizes leads to better generalization on larger problems than using full graphs

Journal : SmallCondensed 10601	Article No : 9327	Pages : 30	MS Code : 9327	Dispatch : 21-4-2022
--------------------------------	-------------------	------------	----------------	----------------------



Fig. 5 Impact of GNN aggregation functions. For larger graphs, aggregators that are agnostic to node degree (MEAN, MAX) are able to outperform theoretically more expressive aggregators

We also experiment with the following normalization schemes: (1) standard BatchNorm which learns mean and variance from training data, as well as (2) BatchNorm with batch statistics; and (3) LayerNorm, which normalizes at the embedding dimension instead of across the batch. Figure 6 indicates that BatchNorm with batch statistics and LayerNorm are able to better account for changing statistics across different graph sizes. Standard BatchNorm generalizes worse than not doing any normalization, thus our other experiments use BatchNorm with batch statistics.

We further dissect the relationship between graph representations and normalization in Appendix D, confirming that poor performance on large graphs can be explained by unstable representations due to the choice of aggregation and normalization schemes. Using MAX aggregators and BatchNorm with batch statistics are temporary hacks to overcome the failure of the current architectural components. Overall, our results suggest that inference beyond training sizes will require the development of expressive GNN mechanisms that are



Fig. 6 Impact of normalization schemes. Modifying BatchNorm to account for changing graph statistics across graph sizes leads to better generalization

Journal : SmallCondensed 10601	Article No : 9327	Pages : 30	MS Code : 9327	Dispatch : 21-4-2022
--------------------------------	-------------------	------------	----------------	----------------------

able to leverage global graph topology [62] while being invariant to distribution shifts interms of node degree and other graph statistics [63].

348 5.4 Which decoder has a better inductive bias for TSP?

Figure 7 compares NAR and AR decoders for identical models. To isolate the impact of the decoder's inductive bias without the inductive bias imposed by GNNs, we also show Transformer encoders on full graphs as well as structure-agnostic MLPs. Within our experimental setup, AR decoders are able to fit the training data as well as generalize significantly better than NAR decoders, indicating that sequential decoding is powerful for TSP even without graph information.

Conversely, NAR architectures are a poor inductive bias as they require significantly more computation to perform competitively to AR decoders. For instance, recent models [17, 18] used more than 30 GNN layers with over 10 Million parameters. We believe that



Fig. 7 Comparing AR and NAR decoders. Sequential AR decoding is a powerful inductive bias for TSP as it enables significantly better generalization, even in the absence of graph structure (MLP encoders)

such overparameterized networks are able to memorize all patterns for small TSP training sizes [64], but the learnt policy is unable to generalize beyond training graph sizes. At the same time, when compared fairly within the same experimental settings, NAR decoders are significantly faster than AR decoders described in Section 3.3 as well as those which reembed the graph at each decoding step [9], see Fig. 8.

363 5.5 How do learning paradigms impact the search phase?

³⁶⁴ Identical models are trained via supervised learning (SL) and reinforcement learning ³⁶⁵ (RL).³ Figure 9 illustrates that, when using greedy decoding during inference, RL mod-³⁶⁶ els perform better on the training size as well as on larger graphs. Conversely, SL models ³⁶⁷ improve over their RL counterparts when performing beam search or sampling.

_{3FL01}³ For RL, we show the greedy rollout baseline. Critic baseline results are available in Appendix E





In Appendix C, we find that the rollout baseline, which encourages better greedy behav-368 iour, leads to the model making very confident predictions about selecting the next node at 369 each decoding step, even out of training size range. In contrast, SL models are trained with 370 teacher forcing, i.e. imitating the optimal solver at each step instead of using their own pre-371 diction. This results in less confident predictions and poor greedy decoding, but makes the 372 373 probability distribution more amenable to beam search and sampling, as shown in Fig. 10. Our results advocate for tighter coupling between the training and inference phase of learn-374 ing-driven TSP solvers, mirroring recent findings in generative models for text [65]. 375

376 5.6 Which learning paradigm scales better?

Our experiments till this point have focused on isolating the impact of various pipeline components on zero-shot generalization under limited computation. At the same time, recent results in natural language processing have highlighted the power of large scale



Fig.9 Comparing solution search settings. Under greedy decoding, RL demonstrates better performance and generalization. Conversely, SL models improve over their RL counterparts when performing beam search or sampling

Journal : SmallCondensed 10601	Article No : 9327	Pages : 30	MS Code : 9327	Dispatch : 21-4-2022
--------------------------------	-------------------	------------	----------------	----------------------



Fig. 10 Impact of increasing beam width. Teacher-forcing during SL leads to poor generalization under greedy decoding, but makes the probability distribution more amenable to beam search

pre-training for effective transfer learning [66]. To better understand the impact of learning paradigms when scaling computation, we double the model parameters (up to 750,000) and train on tens times more data (12.8M samples) for AR architectures. We monitor optimality gap on the training size range (TSP20-50) as well as a larger size (TSP100) vs. the number of training samples.

In Fig. 11, we see that increasing model capacity leads to better learning. Notably, RL 385 models, which train on unique randomly generated samples throughout, are able to keep 386 improving their performance within as well as outside of training size range as they see 387 more samples. On the other hand, SL is bottlenecked by the need for optimal groundtruth 388 solutions: SL models iterate over the same 1.28M unique labelled samples and stop 389 improving at a point. Beyond favorable inductive biases, distributed and sample-efficient 390 RL algorithms [67] may be a key ingredient for learning from and scaling up to larger 391 TSPs beyond tens of nodes. 392

Fig. 11 Scaling computation and parameters for SL and RL-trained models. All models are trained on TSP20-50. We plot optimality gap on 1,280 held-out samples of both TSP50 (performance on training size) and TSP100 (out-of-distribution generalization) under greedy decoding. Note that SL models are less amenable than RL models to greedy search. RL models are able to keep improving their performance within as well as outside of training size range with more data. On the other hand, SL performance is bottlenecked by the need for optimal groundtruth solutions



423

424

425

can generalize to TSPs with up to 1000 nodes. They ensure that the predictions of the GNN 427 encoder generalize from small to large TSP by updating the problem definition (Fig. 2(a), 428 box 1): large problem instances are represented as many smaller sub-graphs which are of 429 the same size as the training graphs for the GNN, and then merge the GNN edge predic-430 tions before performing MCTS. 431

(Fig. 2(a), box 2) while replacing the graph search component of the pipeline (Fig. 2(a),

box 4) with more powerful and flexible algorithms, e.g. Dynamic Programming [31] or

Overall, this line of work suggests that stronger coupling between the design of both 432 the neural and symbolic/search components of models is essential for out-of-distribution 433 generalization. 434

Learning within local search heuristics Recent work has explored an alternative to con-435 structive AR and NAR decoding schemes which involves learning to iteratively improve 436

6 Recent case studies and future work 393

Article No : 9327

Since the initial publication of this work [68], deep learning for routing problems has 394 received considerable attention from the research community [27, 28, 30, 31, 69–73]. In 395 this section, we highlight recent advances, characterize them using the unified pipeline pre-396 sented in Fig. 2, and provide future research directions, with a focus on improving generali-397 398 zation to large-scale and real-world instances.

399 As a reminder, the unified neural combinatorial optimization pipeline consists of: (1) Problem Definition \rightarrow (2) Graph Embedding \rightarrow (3) Solution Decoding \rightarrow (4) Solution 400 Search \rightarrow (5) Policy Learning. 401

Leveraging equivariance and symmetries The autoregressive Attention Model [16] 402 sequentially constructs TSP tours as permutations of cities, but does not consider the 403 underlying symmetries of routing problems. 404

Kwon et al. [27] consider invariance to the starting city in constructive heuristics: They 405 propose to train the Attention Model with a new reinforcement learning algorithm (inno-406 vating on box 5 in Fig. 2(a) which exploits the existence of multiple optimal tour permuta-407 tions. Similarly, Ouyang, Wang, et al. [28] consider invariance with respect to rotations, 408 reflections, and translations (Euclidean symmetry group) of the input cities: They propose 409 a constructive approach similar to Attention Model while ensuring invariance by perform-410 ing data augmentation during the problem definition stage (Fig. 2(a), box 1) and using rela-411 tive coordinates during graph encoding (Fig. 2(a), box 2). Their approach shows particu-412 larly strong results on zero-shot generalization from random instances to the real-world 413 TSPLib benchmark suite. 414

Future work may follow the Geometric Deep :earning blueprint [74] by designing 415 models that respect the symmetries and inductive biases that govern the data. As routing 416 problems are embedded in euclidean coordinates and the routes are cyclical, incorporating 417 these contraints directly into the architectures or learning paradigms may be a principled 418 approach to improving generalization to large-scale instances greater than those seen dur-419 420 ing training.

Improved graph search algorithms Several papers have proposed to improve the one-421 shot non-autoregressive approach of Joshi et al. [17] by retaining the same GNN encoder 422

🖉 Springer

Monte-Carlo Tree Search (MCTS) [30].

Journal : SmallCondensed 10601	4
--------------------------------	---

Article No : 9327

Constraints

(sub-optimal) solutions or learning to perform local search [69–73]. Since deep learning
is used to guide decisions within classical search algorithms (which are designed to work
regardless of problem scale), this approach implicitly leads to better zero-shot generalization to larger problem instances compared to constructive approaches studied in our work.
In particular, NeuroLKH [71] uses GNNs to improve the Lin-Kernighan-Helsgaun algorithm and demonstrates strong zero-shot generalization to TSP with 5000 nodes as well as
across TSPLib instances.

A limitation of this line of work is the need for hand-designed local search heuristics, which may be missing for understudied problems. On the other hand, constructive approaches are comparatively easier to adapt to new problems by enforcing constraints during the solution decoding and search procedure (Fig. 2(a), box 4).

448 Learning Paradigms that promote generalization Future work could look at novel learn-449 ing paradigms which explicitly focus on generalization beyond supervised and reinforce-450 ment learning. For e.g., this work explored zero-shot generalization to larger problems, 451 but the logical next step is to fine-tune the model on a small number of larger problem 452 instances [75]. Thus, it will be interesting to explore fine-tuning/generalization as a meta-453 learning problem, wherein the goal is to train model parameters specifically for fast adapta-454 tion and fine-tuning to new data distributions and problem sizes.

Another interesting direction could explore tackling understudied routing problems with challenging constraints via multi-task pre-training on well-known routing problems such as TSP and CVPR, followed by problem-specific finetuning. Similar to language modelling as a pre-training objective in NLP [66], the goal of pre-training for routing would be to learn generally useful neural network representations that can transfer well to novel routing problems.

461 7 Conclusion

Learning-driven solvers for combinatorial problems such as the Travelling Salesperson Problem have shown promising results for trivially small instances up to a few hundred nodes. However, scaling fully *end-to-end* deep learning approaches to real-world instances is still an open question as training on large graphs is extremely time-consuming and challenging to learn from.

This paper advocates for an alternative to expensive large-scale training: training models efficiently on trivially small TSP and transferring the learnt policy to larger graphs in a *zero-shot* fashion or via fast fine-tuning. Thus, identifying promising inductive biases, architectures and learning paradigms that enable such zero-shot generalization to large and more complex instances is a key concern for tackling real-world combinatorial problems.

We perform the first principled investigation into zero-shot generalization for learning large scale TSP, unifying state-of-the-art architectures and learning paradigms into one experimental pipeline for neural combinatorial optimization. Our findings suggest that key design choices such as GNN layers, normalization schemes, graph sparsification, and learning paradigms need to be explicitly re-designed to consider out-of-distribution generalization. Additionally, we use our unified pipeline to characterize recent advances in deep learning for routing problems and provide new directions to stimulate future research.

479 Appendix A: Additional Context for Fig. 1

Experimental setup In Fig. 1, we illustrate the computational challenges of learning large 480 scale TSP by comparing three identical models trained on 12.8 Million TSP instances via 481 reinforcement learning. Our experimental setup largely follows Section 4. All models use 482 identical configurations: autoregressive decoding and Graph ConvNet encoder with MAX 483 aggregation and LayerNorm. The TSP20-50 model is trained using the greedy rollout base-484 line [16] and the Adam optimizer with batch size 128 and learning rate 1e - 4. Direct train-485 ing, active search and finetuning on TSP200 samples is done using learning rate 1e - 5, as 486 we found larger learning rates to be unstable. During active search and finetuning, we use 487 an exponential moving average baseline, as recommended by Bello et al. [7]. 488

Furthest insertion baseline We characterize 'good' generalization across our experiments by the well-known *furthest insertion* heuristic, which constructively builds a solution/partial tour π' by inserting node *i* between tour nodes $j_1, j_2 \in \pi'$ such that the distance from node *i* to its nearest tour node j_1 is maximized.

We motivate our work by showing that learning from large TSP200 is intractable on 493 university-scale hardware, and that efficient pre-training on trivial TSP20-50 enables mod-494 els to better generalize to TSP200 in a zero-shot manner. Within our computational budget, 495 furthest insertion still outperforms our best models. At the same time, we are not claiming 496 that it is *impossible* to outperform insertion heuristics with current approaches: reinforce-497 ment learning-driven approaches will only continue to improve performance with more 498 computation and training data. We want to use simple non-learnt baselines to motivate the 499 development of better architectures, learning paradigms and evaluation protocols for neural 500 combinatorial optimization. 501

Routing problems and generalization It is worth mentioning why we chose to study TSP in particular. Firstly, TSP has stood the test of time in terms of relevance and continues to serve as an engine of discovery for general purpose techniques in applied mathematics.

TSP and associated routing problems have also emerged as a challenging testbed for 505 learning-driven approaches to combinatorial optimization. Whereas generalization to prob-506 lem instances larger and more complex than those seen in training has at least partially 507 been demonstrated on non-sequential problems such as SAT, MaxCut, MVC [9-11],⁴ the 508 same architectures do not show strong generalization for TSP. For e.g., furthest insertion 509 outperforms or is competitive with state-of-the-art approaches for TSP above tens of nodes, 510 see Fig. D.1.(e, f) from Khalil et al. [9] or Fig. 5 from Kool et al. [16], despite using more 511 computation and data than our controlled study. 512

^{4FL01}⁴ It is worth noting that classical algorithmic and symbolic components such as graph reduction, sophis-^{4FL02}ticated tree search as well as post-hoc local search have been pivotal and complementary to GNNs in ena-^{4FL03}bling such generalization.

	Journal : SmallCondensed 10601	Article No : 9327	Pages : 30	MS Code : 9327	Dispatch : 21-4-2022
--	--------------------------------	-------------------	------------	----------------	----------------------

Table 1Approximate trainingtime (12.8M samples) andinference time (1,280 samples)across TSP sizes and searchsettings for SL and RL-trainedmodels	Graph Size	Training T	ìme	Inference Time		
		SL	RL	GS	BS128	S128
	TSP20	4h 24m	8h 02m	2.62s	7.06s	63.37s
	TSP20-50	9h 49m	15h 47m	—	-	_
	TSP50	16h 11m	40h 29m	7.45s	29.09s	86.48s
	TSP100	68h 34m	108h 30m	19.04s	98.26s	180.30s
	TSP200	_	495h 55m	54.88s	372.09s	479.37s

GS: Greedy search, *BS128*: beam search with width 128, *S128*: sampling 128 solutions. RL training uses the rollout baseline and timing includes the time taken to update the baseline after each 128,000 samples

513 Appendix B: Hardware and Timings

Fairly timing research code can be difficult due to differences in libraries used, hardware configurations and programmer skill. In Table 1, we report approximate total training time and inference time across TSP sizes for the model setup described in Section 4. All experiments were implemented in PyTorch and run on an Intel Xeon CPU E5-2690 v4 server and four Nvidia 1080Ti GPUs. Four experiments were run on the server at any given time (each using a single GPU). Training time may vary based on server load, thus we report the lowest training time across several runs in Table 1.

521 Appendix C: Learning paradigms and amenity to search

Figure 10 demonstrate that SL models are more amenable to beam search and sampling, 522 but are outperformed by RL-rollout models under greedy search. In Fig. 12, we investigate 523 the impact of learning paradigms on probability distributions by plotting histograms of the 524 probabilities of greedy selections during inference across TSP sizes for identical models 525 trained with SL and RL. We find that the rollout baseline, which encourages better greedy 526 behaviour, leads to the model making very confident predictions about selecting the next 527 node at each decoding step, even beyond training size range. In contrast, SL models are 528 trained with teacher forcing, i.e. imitating the optimal solver at each step instead of using 529 their own prediction. This results in less confident predictions and poor greedy decod-530 ing, but makes the probability distribution more amenable to beam search and sampling 531 532 techniques.

We understand this phenomenon as follows: More confident predictions (Fig. 12b) do not automatically imply better solutions. However, sampling repeatedly or maintaining the top-*b* most probable solutions from such distributions is likely to contain very similar tours. On the other hand, less sharp distributions (Fig. 12a) are likely to yield more diverse tours with increasing *b*. This may result in comparatively better optimality gap, especially for TSP sizes larger than those seen in training.





539 Appendix D: Visualizing node and graph embedding spaces

Our results in Section 5.3 suggest that inference beyond training sizes requires the devel-540 opment of GNN architectures and normalization layers that are both expressive as well 541 as invariant to distribution shifts. We explore how node and graph embeddings for TSP 542 graphs evolve across training distribution (TSP20-50) and beyond (up to TSP200) through 543 visualizing the statistics of the embedding spaces. Intuitively, constructing TSP tours 544 involves decisions which are not just locally optimal, but also optimal w.r.t some global 545 graph structure. Thus, node embeddings represent *local* information while graph embed-546 dings, which are conventionally computed as the mean of node embeddings, provide global 547 structural information. 548

We utilize distribution plots to study the variation in embedding statistics⁵ of three identical models: (1) **GNN-Max**, which represents our best model configuration from Section 5: autoregressive decoding, Graph ConvNet encoder with MAx aggregation and BatchNorm with batch statistics; (2) **GNN-Sum**, which uses SUM aggregation for the Graph ConvNet and shows comparatively poor generalization beyond training size, see Fig. 5; and (3) **GNN-Max + learnt BN**, which uses standard BatchNorm, i.e. learns statistics from the training data, and also shows comparatively poor generalization, see Fig. 6.

We draw upon work in learning embeddings for computer vision [77] to characterize embedding spaces across TSP sizes according to: (1) **magnitudes**, denoted by ℓ_2 norms, indicating whether embeddings are shrinking to one magnitude or expanding outwards as TSP size increases; and (2) **pair-wise distances**, which tells us how well-separated the embedding are, or whether they are pulled apart/towards each other as TSP size increases.

^{5FL01}⁵ Distribution plots show 0, 5, 50, 95, and 100-percentiles for embedding statistics at various TSP sizes, ^{5FL02}thus visualizing how the statistics changes with problem scale (implemented via TensorBoard [76]).

Journal : SmallCondensed 10601	Article No : 9327	Pages : 30	MS Code : 9327	Dispatch : 21-4-2022



Fig. 13 Distribution plots of node embedding ℓ_2 norms (y-axis) across TSP sizes (x-axis)



Fig. 14 Distribution plots of node embedding pair-wise distances (y-axis) across TSP sizes (x-axis)

Node embedding space In Figs. 13 and 14, we see that *GNN-Max* leads to the most sta-561 ble node embedding norms and pair-wise distances (which are calculated at an intra-graph 562 level) across TSP sizes. On the other hand, GNN-Sum and GNN-Max + learnt BN lead to 563 fluctuating and monotonically increasing embedding norms as size increases, e.g. compare 564 Fig. 13b and c. Clearly, maintaining similar distributions for node embeddings across graph 565 sizes indicates that the GNN is building meaningful representations of local structure, or, 566 at the very least, does not break down for large graphs. This enables better generalization, 567 as the decoder has lower chances of encountering embeddings which are statistically differ-568 ent than those seen during training. 569

Graph embedding space Figures 15 and 16 indicate that the graph embedding space is shrinking towards a single magnitude and moving closer as graph size increases. Interestingly, with standard BatchNorm, the graph embedding magnitude monotonically increases with graph size to ranges beyond those for training graphs. On the other hand, using batch statistics for BatchNorm, as done in *GNN-Max* and *GNN-Sum*, leads to graph embedding magnitudes converging to a single value which is within the range of values for training graphs, thus enabling better generalization. E.g. compare Fig. 15b and c.



Fig. 15 Distribution plots of graph embedding ℓ_2 norms (y-axis) across TSP sizes (x-axis)



Fig. 16 Distribution plots of graph embedding pair-wise distances (y-axis) across TSP sizes (x-axis)

We can further visualize this phenomenon through 2D Principal Component Analysis (PCA) plots of graph embedding spaces for *GNN-Max* and *GNN-Max* + *learnt BN* models, see Fig. 17a and b. In both cases, the graph embeddings at larger sizes have very similar magnitudes and are extremely close to each other, indicating that the model is unable to differentiate among different graphs. Thus, decoders currently lack good global structural context. Investigating better graph embeddings through pooling methods [78] could be an interesting approach towards representing global graph structure beyond training sizes.

584 Appendix E: Extra Results

NAR decoders and aggregation functions In Section 5, we found that AR decoding pro-585 vides a powerful sequential inductive bias for TSP and is able to generalizes well with 586 both GNNs as well as structure-agnostic encoder architectures. This result may lead one 587 to question the need for GNNs, altogether. Interestingly, Fig. 18 illustrates a different trend 588 for NAR architectures: GNN encoders generalize better than both Transformers and MLPs, 589 indicating that leveraging graph structure is essential in the absence of the sequential 590 inductive bias. (It is worth noting that, overall, all models with NAR decoders generalize 591 poorly compared to AR architectures for our experimental setup.) 592

Critic baseline Figure 19 illustrates that, for identical models, the critic baseline [7, 19] is unable to match the performance of the rollout baseline [16] under both greedy and beam search settings. We did not explore tuning learning rates and hyperparameters for the critic network, opting to use the same settings as those for the actor. In general, getting

L	A	D	M6 G 1 0227	Direct 1 - 21 4 2022
Journal : SmallCondensed 10601	Article No : 9327	Pages : 30	MS Code : 9327	Dispatch : 21-4-2022



(a) GNN-Max

(b) GNN-Max + learnt BN

Fig. 17 2D PCA of graph embedding spaces. Colors represent TSP instance sizes, e.g. orange: TSP10, teal: TSP20, pink: TSP50, dark grey: TSP200



597 actor-critic methods to work seems to require more parameter tuning than the rollout 598 baseline.

Scaling computation for AR and NAR architectures In Figures 20 and 21, we present 599 extended results for Section 5.6, where we scale model parameters and data. We observe 600 that using larger models (up to 1.5 Million parameters) enables fitting the training dataset 601 better. The impact of larger models is especially evident for NAR architectures. As previ-602 ously noted, recent NAR-based models [17, 18] used more than 30 layers with over 10 603 Million parameters to outperform AR architectures on fixed TSP sizes. We believe that 604 such overparameterized networks are able to memorize all patterns for small TSP training 605 sizes [64], but the learnt policy is unable to generalize beyond training graph sizes as NAR 606 decoding does not provide a useful inductive bias for TSP. 607

Journal : SmallCondensed 10601	Article No : 9327	Pages : 30	MS Code : 9327	Dispatch : 21-4-2022
--------------------------------	-------------------	------------	----------------	----------------------



Fig. 20 Scaling computation and model parameters for AR decoder

Journal : SmallCondensed 10601	Article No : 9327	Pages : 30	MS Code : 9327	Dispatch : 21-4-2022
--------------------------------	-------------------	------------	----------------	----------------------



Fig. 21 Scaling computation and model parameters for NAR decoder

Journal	:	SmallCondensed	10601	
---------	---	----------------	-------	--

608 Appendix F: Visualizing model predictions

As a final note, we present a visualization tool for generating model predictions and heatmaps of TSP instances, see Figures 22, 23. We advocate for the development of more principled approaches to neural combinatorial optimization, e.g. along with model predictions, visualizing the reduce costs for each edge (cheaply obtained using the Gurobi solver [79]) may help debug and improve learning-driven approaches in the future. Using reduce costs as supervision signals could also be an inexpensive alternative to running optimal solvers to create large labelled datasets.



Fig. 22 Prediction visualization for TSP20



Fig. 23 Prediction visualization for TSP50

Journal : SmallCondensed 10601	Article No : 9327	Pages : 30	MS Code : 9327	Dispatch : 21-4-2022
--------------------------------	-------------------	------------	----------------	----------------------

616 Acknowledgements We would like to thank R. Anand, X. Bresson, V. Dwivedi, A. Ferber, E. Khalil, 617 W. Kool, R. Levie, A. Prouvost, P. Veličković and the anonymous reviewers for helpful comments and 618 discussions.

619 **Declarations**

620 **Conflict of Interests** None.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, 621 which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long 622 as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Com-623 mons licence, and indicate if changes were made. The images or other third party material in this article 624 are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the 625 material. If material is not included in the article's Creative Commons licence and your intended use is not 626 permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly 627 from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/. 628

629 References

- Lenstra, J.K., & Kan, A.R. (1975). Some simple applications of the travelling salesman problem. *Journal of the Operational Research Society*.
- Applegate, D.L., Bixby, R.E., Chvatal, V., & Cook, W.J. (2006). The traveling salesman problem: A
 computational study.
- 634 3. Applegate, D., Bixby, R., Chvatal, V., & Cook, W. (2006). Concorde TSP solver.
- Senior, A.W., Evans, R., Jumper, J., Kirkpatrick, J., Sifre, L., Green, T., Qin, C., žídek, A., Nelson,
 A.W., Bridgland, A., & et al. (2020). Improved protein structure prediction using potentials from deep
 learning. *Nature*.
- 5. Mirhoseini, A., Goldie, A., Yazgan, M., Jiang, J.W., Songhori, E., Wang, S., Lee, Y.-J., Johnson, E.,
 Pathak, O., Nazi, A., & et al. (2021). A graph placement methodology for fast chip design. *Nature*.
- 640 6. Vinyals, O., Fortunato, M., & Jaitly, N. (2015). Pointer networks. In NeurIPS.
- 7. Bello, I., Pham, H., Le, Q.V., Norouzi, M., & Bengio, S. (2017). Neural combinatorial optimization with reinforcement learning. In *ICLR*.
- 8. Bengio, Y., Lodi, A., & Prouvost, A. (2020). Machine learning for combinatorial optimization: a methodological tour d'horizon. *European Journal of Operational Research*.
- 645 9. Khalil, E., Dai, H., Zhang, Y., Dilkina, B., & Song, L. (2017). Learning combinatorial optimization
 646 algorithms over graphs. In *NeurIPS*.
- 647 10. Selsam, D., Lamm, M., Bünz, B., Liang, P., de Moura, L., & Dill, D.L. (2019). Learning a sat solver
 648 from single-bit supervision. In *ICLR*.
- Li, Z., Chen, Q., & Koltun, V. (2018). Combinatorial optimization with graph convolutional networks
 and guided tree search. In *NeurIPS*.
- Kipf, T.N., & Welling, M. (2017). Semi-supervised classification with graph convolutional networks.
 In *ICLR*.
- 653 13. Gilmer, J., Schoenholz, S.S., Riley, P.F., Vinyals, O., & Dahl, G.E. (2017). Neural message passing for 654 quantum chemistry. In *ICML*.
- 14. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., & Bengio, Y. (2018). Graph attention
 networks. ICLR.
- Battaglia, P.W., Hamrick, J.B., Bapst, V., Sanchez-Gonzalez, A., Zambaldi, V., Malinowski, M., Tacchetti, A., Raposo, D., Santoro, A., Faulkner, R., & et al. (2018). Relational inductive biases, deep
 learning, and graph networks. arXiv preprint.
- 660 16. Kool, W., van Hoof, H., & Welling, M. (2019). Attention, learn to solve routing problems!. In ICLR.
- I7. Joshi, C.K., Laurent, T., & Bresson, X. (2019). An efficient graph convolutional network technique for
 the travelling salesman problem. arXiv preprint.
- 18. Nowak, A., Villar, S., Bandeira, A.S., & Bruna, J. (2017). A note on learning algorithms for quadratic
 assignment with graph neural networks. arXiv preprint.
- 19. Deudon, M., Cournut, P., Lacoste, A., Adulyasak, Y., & Rousseau, L.-M. (2018). Learning heuristics
 for the TSP by policy gradient. In *CPAIOR*.

Journal: SmanCondensed 10001 Article No : 9327 Pages : 50 Mis Code : 9327 Dispatch : 21-4	Journal : SmallCondensed 10601	Article No : 9327	Pages : 30	MS Code : 9327	Dispatch : 21-4-2
---	--------------------------------	-------------------	------------	----------------	-------------------

		Constraints
	20	Gasse M. Chátelat D. Ferroni N. Charlin I. & Lodi A. (2019). Exact combinatorial optimization
667	20.	with granh convolutional neural networks. In <i>NeurIPS</i>
668	21	Cappart O Goutierre E Bergman D & Rousseau I M (2010) Improving optimization bounds
669	21.	using machine learning: Decision diagrams meet deen reinforcement learning. In AAAI
670	22	Chalumeau F Coulon I Cannart O & Rousseau I -M (2021) Seanearl: A constraint program-
670	22.	ming solver guided by reinforcement learning. In CPAIOR
672	23	Wilder B. Dilkina B. & Tambe M (2019) Melding the data-decisions pipeline: Decision-focused
674	20.	learning for combinatorial optimization. In AAAI
675	24	Ferber A Wilder B Dilking B & Tambe M (2020) MIPaal : Mixed integer program as a layer. In
676	2	AAAI.
677	25.	Sutskever, I., Vinvals, O., & Le, O.V. (2014). Sequence to sequence learning with neural networks. In
678		NeurIPS.
679	26.	Ma, Q., Ge, S., He, D., Thaker, D., & Drori, I. (2020). Combinatorial optimization by graph pointer
680		networks and hierarchical reinforcement learning. In AAAI workshop on deep learning on graphs.
681	27.	Kwon, YD., Choo, J., Kim, B., Yoon, I., Gwon, Y., & Min, S. (2020). Pomo: Policy optimization with
682		multiple optima for reinforcement learning. In NeurIPS.
683	28.	Ouyang, W., Wang, Y., Weng, P., & Han, S. (2021). Generalization in deep rl for tsp problems via
684		equivariance and local search. arXiv preprint.
685	29.	Nowak, A., Folqué, D., & Estrach, J.B. (2018). Divide and conquer networks. In ICLR.
686	30.	Fu, ZH., Qiu, KB., & Zha, H. (2021). Generalize a small pre-trained model to arbitrarily large TSP
687		instances. In AAAI.
688	31.	Kool, W., van Hoof, H., Gromicho, J., & Welling, M. (2021). Deep policy dynamic programming for
689		vehicle routing problems. arXiv preprint.
690	32.	Joshi, C.K., Laurent, T., & Bresson, X. (2019). On learning paradigms for the travelling salesman
691		problem. NeurIPS Graph Representation Learning Workshop.
692	33.	Nazari, M., Oroojlooy, A., Snyder, L., & Takác, M. (2018). Reinforcement learning for solving the
693	24	vehicle routing problem. In <i>NeurIPS</i> .
694	34.	Chen, X., & Tian, Y. (2019). Learning to perform local rewriting for combinatorial optimization. In
695	25	NeurIPS.
696	35. 26	YOICU, E., & POCZOS, B. (2019). Learning local search neuristics for boolean satisfiability. In <i>NeurIPS</i> .
697	30. 27	Future, J., Patwary, M., & Diamos, G. (2019). Coloring olg graphs with alphagozeto. arXiv preprint.
698	57.	Sato, K., Tamada, M., & Kasmina, H. (2019). Approximation ratios of graph neural networks for com-
699	38	Cappart O. Chátalat D. Khalil F. Lodi A. Morris C. & Valičković P. (2021). Combinatorial anti-
700	50.	mization and reasoning with graph neural networks. In <i>IICAI</i>
701	30	Veličković P Ving R Padovano M Hadsell R & Blundell C (2020) Neural execution of granh
702	57.	algorithms In <i>JCLR</i>
703	40	Veličković, P., & Blundell, C. (2021). Neural algorithmic reasoning. <i>Patterns</i>
704	41	Corso, G., Cavalleri, L., Beani, D., Liò, P., & Veličković, P. (2020). Principal neighbourhood aggrega-
706		tion for graph nets. In <i>NeurIPS</i>

- Xu, K., Li, J., Zhang, M., Du, S.S., Kawarabayashi, K.-i., & Jegelka, S. (2019). What can neural net-42. 707 works reason about?. In ICLR. 708
- 43. Xu, K., Li, J., Zhang, M., Du, S.S., Kawarabayashi, K.-i., & Jegelka, S. (2020). How neural networks 709 710 extrapolate: From feedforward to graph neural networks. In ICLR.
- Gómez-Bombarelli, R., Wei, J.N., Duvenaud, D., Hernández-Lobato, J.M., Sánchez-Lengeling, 44. 711 B., Sheberla, D., Aguilera-Iparraguirre, J., Hirzel, T.D., Adams, R.P., & Aspuru-Guzik, A. (2018). 712 Automatic chemical design using a data-driven continuous representation of molecules. ACS central 713 science. 714
- Mao, H., Schwarzkopf, M., Venkatakrishnan, S.B., Meng, Z., & Alizadeh, M. (2019). Learning sched-45. 715 uling algorithms for data processing clusters. In ACM special interest group on data communication. 716
- Paliwal, A., Gimeno, F., Nair, V., Li, Y., Lubin, M., Kohli, P., & Vinyals, O. (2019). Regal: Transfer 46. 717 learning for fast optimization of computation graphs. arXiv preprint. 718
- Mirhoseini, A., Pham, H., Le, Q.V., Steiner, B., Larsen, R., Zhou, Y., Kumar, N., Norouzi, M., Bengio, 47. 719 S., & Dean, J. (2017). Device placement optimization with reinforcement learning. In ICML. 720
- 48. Zhou, Y., Roy, S., Abdolrashidi, A., Wong, D., Ma, P.C., Xu, Q., Zhong, M., Liu, H., Goldie, A., 721 Mirhoseini, A., & et al. (2019). Gdp: Generalized device placement for dataflow graphs. arXiv 722 preprint. 723
- 49. Bresson, X., & Laurent, T. (2019). A two-step graph convolutional decoder for molecule generation. In 724 NeurIPS workshop on machine learning and the physical sciences. 725

Journal : SmallCondensed	10601
--------------------------	-------

MS Code : 9327

Constraints

726	50.	Jin, W., Barzilay, R., & Jaakkola, T. (2018). Junction tree variational autoencoder for molecular graph
727	51	generation. In <i>ICML</i> .
728	51.	You, J., Liu, B., Ying, Z., Pande, V., & Leskovec, J. (2018). Graph convolutional policy network for goal directed molecular graph generation. In <i>NeurIPS</i>
729	52	Bresson X & Laurent T (2018) An experimental study of neural networks for variable graphs. In
731	02.	ICLR Workshop.
732	53.	Ioffe, S., & Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing
733		internal covariate shift. arXiv preprint.
734	54.	Ba, J.L., Kiros, J.R., & Hinton, G.E. (2016). Layer normalization. arXiv preprint.
735	55.	Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., & Polosukhin,
736		I. (2017). Attention is all you need. In <i>NeurIPS</i> .
737	56.	Joshi, C. (2020). Transformers are graph neural networks. The Gradient.
738	57.	Dwivedi, V.P., Joshi, C.K., Laurent, T., Bengio, Y., & Bresson, X. (2020). Benchmarking graph neural
739	50	networks. arXiv preprint.
740	58.	François, A., Cappari, Q., & Rousseau, LIVI. (2019). How to evaluate machine learning approaches
741	59	Williams R L & Zinser D (1989) A learning algorithm for continually running fully recurrent neu-
742	57.	ral networks. <i>Neural Computation</i> 1(2), 270–280
743	60.	Williams, R.J. (1992). Simple statistical gradient-following algorithms for connectionist reinforcement
745		learning. Machine Learning.
746	61.	Xu, K., Hu, W., Leskovec, J., & Jegelka, S. (2019). How powerful are graph neural networks?. In
747		ICLR.
748	62.	Garg, V.K., Jegelka, S., & Jaakkola, T. (2020). Generalization and representational limits of graph neu-
749		ral networks. In <i>ICML</i> .
750	63.	Levie, R., Bronstein, M.M., & Kutyniok, G. (2019). Transferability of spectral graph convolutional
751	64	neural networks. arXiv preprint.
752	64.	Zhang, C., Bengio, S., Hardt, M., Recht, B., & Vinyais, O. (2017). Understanding deep learning
753	65	Holtzman A Buys I Du I Forbes M & Choi V (2020) The curious case of neural text decen
754	05.	eration In ICLR
755	66.	Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., & Liu, P.J.
757		(2020). Exploring the limits of transfer learning with a unified text-to-text transformer. JMLR.
758	67.	Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). Proximal policy optimiza-
759		tion algorithms. arXiv preprint.
760	68.	Joshi, C.K., Cappart, Q., Rousseau, LM., & Laurent, T. (2021). Learning tsp requires rethinking gen-
761		eralization. In International conference on principles and practice of constraint programming.
762	69.	Wu, Y., Song, W., Cao, Z., Zhang, J., & Lim, A. (2021). Learning improvement heuristics for solving
763	70	routing problem. <i>IEEE Transactions on Neural Networks and Learning Systems</i> .
764	70.	da Costa, P.R.d.O., Rhuggenaath, J., Zhang, Y., & Akcay, A. (2020). Learning 2-opt heuristics for the
765	71	Tavening salesman problem via deep reminiscement learning. In Asian conference on machine learning. Xin L. Song W. Cao, 7, & Zhang L (2021) Neurolkh: Combining deep learning model with lin
766	/1.	kernighan helsgaun heuristic for solving the traveling salesman problem. In <i>NeurIPS</i>
769	72.	Ma. Y., Li, J., Cao, Z., Song, W., Zhang, L., Chen, Z., & Tang, J. (2021). Learning to iteratively solve
769		routing problems with dual-aspect collaborative transformer. In <i>NeurIPS</i> .
770	73.	Hudson, B., Li, Q., Malencia, M., & Prorok, A. (2021). Graph neural network guided local search for
771		the traveling salesperson problem. arXiv preprint.
772	74.	Bronstein, M.M., Bruna, J., Cohen, T., & Veličković, P. (2021). Geometric deep learning: Grids,
773		groups, graphs, geodesics, and gauges. arXiv preprint.
774	75.	Hottung, A., Kwon, YD., & Tierney, K. (2021). Efficient active search for combinatorial optimization
775	70	problems. arXiv preprint.
776	/6.	Adadi, M., Agarwai, A., Barnam, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G.S., Davis, A., Dean, J. Davin, M., & et al. (2016). Tangarflow: Large scale machine learning on heterogeneous distributed
777		J., Devin, M., & et al. (2010). Tensornow. Large-scale machine rearning on neterogeneous distributed
770	77	Hermans, A., Bever, L., & Leibe, B. (2017). In defense of the triplet loss for person re-identification.
780		arXiv preprint.
781	78.	Ying, Z., You, J., Morris, C., Ren, X., Hamilton, W., & Leskovec, J. (2018). Hierarchical graph repre-
782		sentation learning with differentiable pooling. In NeurIPS.
783	79.	Inc, G.O. (2015). Gurobi optimizer reference manual. URL http://www.gurobi.com.

Journal : SmallCondensed 10601	Article No : 9327	Pages : 30	MS Code : 9327	Dispatch : 21-4-2022

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

786

weeken

Journal:	10601
Article:	9327

Author Query Form

Please ensure you fill out your response to the queries raised below and return this form along with your corrections

Dear Author

During the process of typesetting your article, the following queries have arisen. Please check your typeset proof carefully against the queries listed below and mark the necessary changes either directly on the proof/online grid or in the 'Author's response' area provided below

Query	Details Required	Author's Response
AQ1	Footnotes are not allowed in abstract, thus we captured it as "Article note". Please check if okay.	
AQ2	(City) has been provided in affiliation 1-3, please check if it is correct.	
AQ3	Please check all Equations if captured and presented correctly	