

FORM FOLLOWS FUNCTION: TEXT-TO-TEXT CONDITIONAL GRAPH GENERATION BASED ON FUNCTIONAL REQUIREMENTS

Anonymous authors

Paper under double-blind review

ABSTRACT

This work focuses on the novel problem setting of generating graphs conditioned on a description of the graph’s functional requirements in a downstream task. We pose the problem as a text-to-text generation problem and focus on the approach of fine-tuning a pretrained large language model (LLM) to generate graphs. We propose an inductive bias which incorporates information about the structure of the graph into the LLM’s generation process by incorporating message passing layers into an LLM’s architecture. To evaluate our proposed method, we design a novel set of experiments using publicly available and widely studied molecule and knowledge graph data sets. Results suggest our proposed approach generates graphs which more closely meet the requested functional requirements, outperforming baselines developed on similar tasks by a statistically significant margin.

1 INTRODUCTION

Many concepts can be described by graphs; including molecules (20), abstract syntax trees (19), knowledge graphs (27), and project schedules (18). Each of these concepts is used in downstream tasks where graph structure has a direct impact on task performance. For example, some molecules can be used as medicine for a disease while others cannot and this is partially determined by the molecules’ graphs. It would be useful if we could describe the functional requirements of a graph using natural language and query a model to conditionally generate a graph which meets these requirements. For example, a model using the prompt "generate a molecule with 47 valency electrons" would generate a valid molecule graph with 47 valency electrons. With such models we could speed up and improve the processes of drug discovery, software development, project management as well as many other applications of graphs.

In this work, we focus on the problem of generating graphs where node and edge features are strings of text and term this type of graph a *text graph*. Text graphs are a fairly flexible data format and can be used in most applications of graph-structured data including those listed above. To the best of our knowledge, text graph generation conditioned on text has only been studied in the fields of knowledge graph generation (8; 12; 14; 26; 27; 38) and explanation graph generation (33). In both setups the conditional text explicitly describes the graph. These explicit descriptions give an imperative "recipe" with step-by-step instructions of how to generate the graph. Contrary to giving explicit imperative descriptions, this work focuses on the case where the conditional text is a *functional* description of the graph, specifically for a downstream task. Additionally, in the case of imperative conditional text, there is an injective mapping between conditional text and the graph that should be generated, which means that methods can be evaluated by comparing the generated graph to the ground truth graph in terms of its structure. By contrast, many graphs might correspond to the same functional description and hence a new experimental design is required to study the problem.

We propose solving this problem by fine tuning a large language model (LLM) to generate a serialized text description of a graph. This is motivated by the strong performance pre-trained LLMs have shown when fine-tuned to perform a specific task (6; 17; 34), including the task of text-to-serialized graph generation (8). However, prior work on serialized graph generation does not generalize to all domains containing text graph structured data. To further this approach, we propose a serialization method which is expressive enough to reversibly serialize any graph and a method to incorporate

graph structure into the LLM’s generation process motivated by the observation that the structures of generated graphs have a direct impact on their functional properties and consequently on model performance. While LLMs are probably expressive enough to learn to incorporate graph structure into their generation process, it is more efficient to provide the model with the ability to do so.

However, it is also a non-trivial challenge because modern LLMs perform autoregressive sampling to generate text (4; 5; 6; 17; 28; 31; 34). This involves a multi-step sequential process, where at each step a forward pass is performed and then a new token is sampled. A token is an atomic unit of text like ‘dog’, ‘ ’, or ‘ch’. The generation process exhibits high computational complexity and cost due to the requirement of performing a forward model pass for each generated token. Consequently, LLMs are typically trained to generate text without actually performing generation at training time. Only a single forward pass is performed at each training step without sampling and the entire sequence (conditional plus desired text) is fed into the LLM for that forward pass. As such, an LLM has access to the full sequence at training time, while at generation time it only has access to the conditional text and the part of the sequence that has already been generated. This means that during training an LLM could learn to rely on information that it will not have access to at generation time. This issue is overcome by using causal masking, which guarantees the LLM cannot pass information backwards in the token sequence ensuring that generation conditions are simulated at training time.

In Section 3, we propose a method of incorporating graph structure into an LLM’s generation process which passes information between tokens in a sequence. Our key contribution is demonstrating how to do so without passing information backwards in the sequence. **Specifically, we propose to extend LLMs to process and generate sequences of text and graphs. For this we provide the LLM with the ability to deserialize graphs incorporated into input token sequences and introduce message passing layers into the LLM to calculate representations of graph structure in a manner that is conducive to autoregressive generation.**

In this work, we specifically focus on the problem of generating text graphs from functional requirements within the range of those seen at training time (interpolation), instead of extrapolating outside the training set’s domain, which we leave as future work. We propose a novel experiment design to evaluate methods in this new problem setting using the publicly available data sets WebNLG+ 2020 (10) and PCQM4M (20). Results suggest that our proposed model outperforms previous work on text graph generation conditioned on text by a statistically significant margin. And that our proposed approach for incorporating graph structure into the language model’s generation process leads to models which generate examples which on average more closely meet their conditional functional requirements. All code used to perform experiments is publicly available at ...

2 PRELIMINARIES

A text graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ is composed of a set of N nodes $v_i \in \mathcal{V}$ each with an associated text string describing the node and a set of M directed edges $(v_i, v_j) \in \mathcal{E}$ each with an associated text string describing the edge. Each text graph \mathcal{G} is associated with text \mathcal{D}_f containing a functional description of \mathcal{G} . Our goal is to train a model to accurately predict $p_{\theta}(\mathcal{G}|\mathcal{D}_f)$ with parameters θ , describing the distribution over text graphs \mathcal{G} conditioned on a specific functional description \mathcal{D}_f . Importantly, this distribution may be multimodal as multiple graphs may have the same functional properties.

This work focuses on the case where the parameterized model $p_{\theta}(\cdot)$ is a language model that generates a serialized text description of the graph. Hence, if $\mathcal{D}_{\mathcal{G}}$ is the serialized graph, then the language model predicts $p_{\theta}(\mathcal{D}_{\mathcal{G}}|\mathcal{D}_f)$. This requires an injective serialization function $g : \mathcal{G} \rightarrow \mathcal{D}_{\mathcal{G}}$ with a known inverse mapping $g^{-1} : \mathcal{D}_{\mathcal{G}} \rightarrow \mathcal{G}$ from the serialized description $\mathcal{D}_{\mathcal{G}}$ back to the graph \mathcal{G} such that $\mathcal{G} = g^{-1}(g(\mathcal{G}))$. $g(\cdot)$ is used at training time to transform graphs \mathcal{G} into serialized graphs $\mathcal{D}_{\mathcal{G}}$ and the deserialization function is used at generation time to recover the graph \mathcal{G} from the serialized graph $\mathcal{D}_{\mathcal{G}}$ (see Figure 1 top row). With this approach, we pose the problem of generating graphs conditioned on functional descriptions as a text conditioned on text generation problem and can consequently use LLMs to conditionally generate text graphs based on functional descriptions.

Prior state-of-the-art work on serialized graph generation uses a bag-of-edges approach (8; 33), which describes each edge using a string with special tokens for delimiting graph elements and then lists out the edges in a predefined order. For example, the popular smiles molecule serialisation algorithm uses depth-first search traversal (39) to order chemical bonds (edges) in a

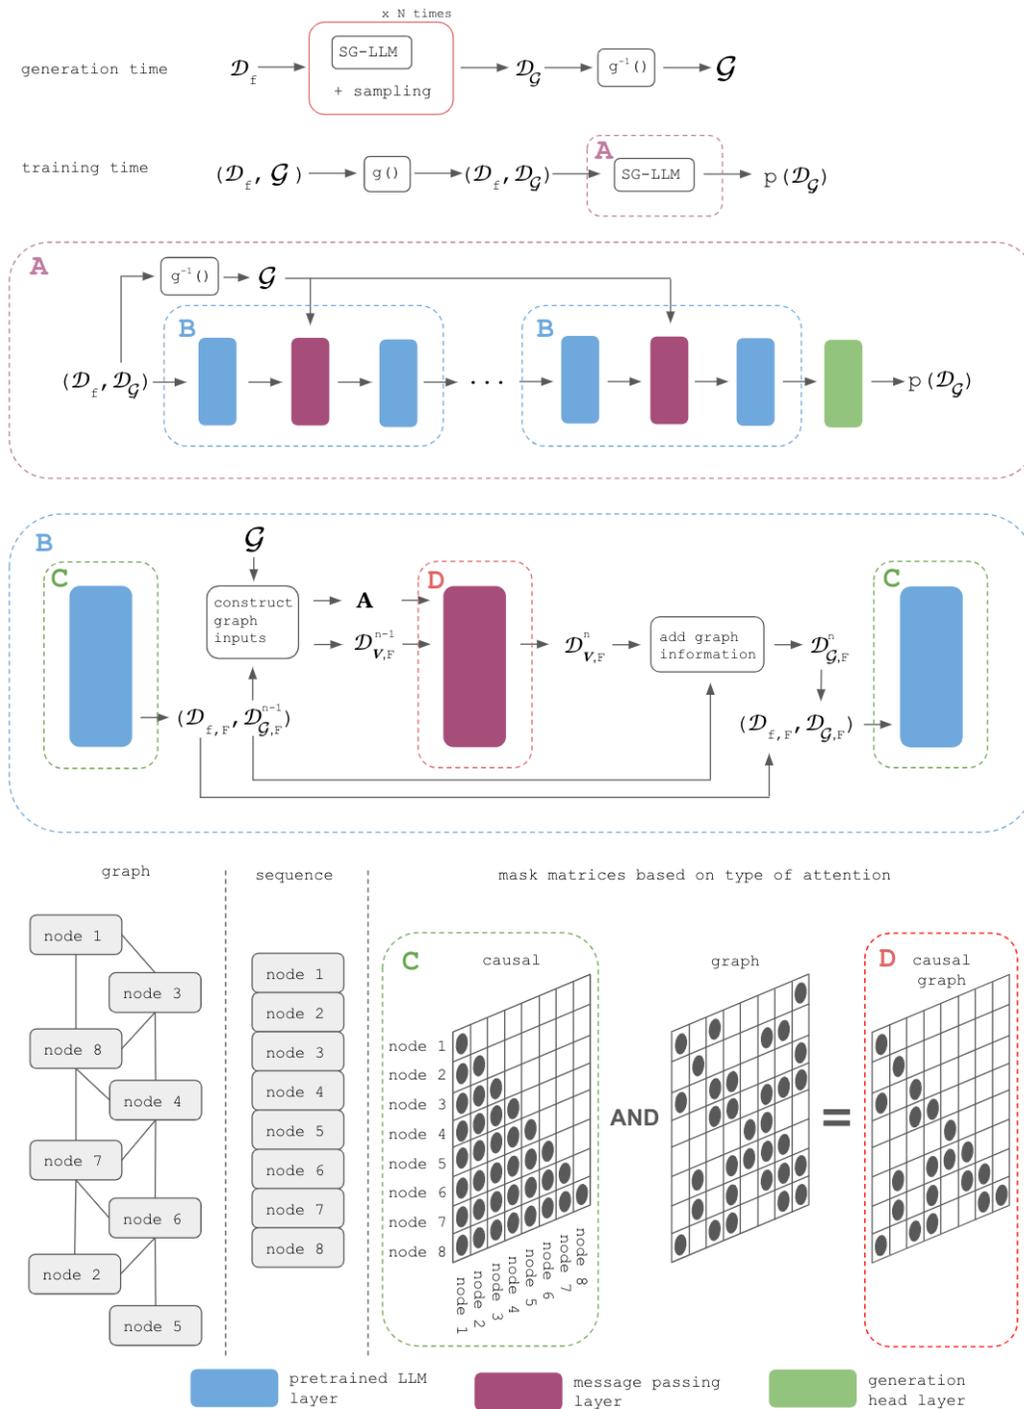


Figure 1: At generation time (top) the input to the model is the graph’s functional description \mathcal{D}_f and the output is a serialised description of the graph \mathcal{D}_g . During training time (below top) the input is a functional description, serialized graph pair $(\mathcal{D}_f, \mathcal{D}_g)$; the output is the probability $p(\mathcal{D}_g)$. Graph serialization method $g(\cdot)$ is used to serialize input graphs before training, while the deserialization method $g^{-1}(\cdot)$ is used to deserialize the generated serialized graph \mathcal{D}_g as well as to help perform message passing within the LLM. Block A describes the proposed architecture interleaving message passing layers between LLM layers. Block B depicts how information is passed between LLM and MP layers. The bottom visualization depicts the masked matrices used in the types of attention performed within LLM and message passing (MP) layers.

molecule. Then each edge is described with the syntax `<PN>predecessor node feature string<E>edge feature string<SN>successor node feature string`, where `<PN>`, `<E>`, `<SN>` are special tokens used to delimit the different components of the edge.

Using LLMs, a description \mathcal{D} is represented in three ways: 1) as a string (text), 2) as a sequence of token indicators, and 3) as a sequence of feature vectors / embeddings. To distinguish between these three representations, we use \mathcal{D}_S to denote the string representation, \mathcal{D}_T to denote the sequence of token indicators and \mathcal{D}_F to denote the sequence of feature vectors. For the serialized graph \mathcal{D}_G , its string representation is denoted by $\mathcal{D}_{G,S}$, its sequence of token indicators representation by $\mathcal{D}_{G,T}$ and its sequence of feature vector representation by $\mathcal{D}_{G,F}$.

3 METHOD

The datasets LLMs are pre-trained on do not contain much, if any, data on generating serialized graphs (4; 6; 31; 34). Consequently, it is unlikely they could generate text graphs without further training. We provide evidence supporting this hypothesis in experiments where we evaluate a pretrained LLM without further fine-tuning. As shown in the results in Appendix D, an LLM without fine-tuning did not generate a single valid serialized graph. Hence, we propose fine-tuning an LLM on a training and validation set of (functional description, graph) pairs to generate serialized graphs conditioned on functional descriptions. While LLMs are probably expressive enough to learn to incorporate graph structure into their generation process, it is more efficient to provide the model with the ability to do so. State-of-the-art models for performing inference on graph structured data use message passing (MP) layers (15; 35). As such, we propose interleaving MP layers between language modelling layers in an LLM (Figure 1 block A). Any message passing layer can be used with our proposed approach ¹.

3.1 FINE-TUNING OBJECTIVE

For fine-tuning, we propose minimizing the negative log-likelihood of serialized graph sequences \mathcal{D}_G as described in equation (1):

$$Loss = \frac{1}{n_{batch}} \sum_{i=1}^{n_{batch}} \left(-\log p_{\theta}(d_{G,T}^{1,i} | \mathcal{D}_f^i) + \sum_{j=2}^{n_{seq}^i} -\log p_{\theta}(d_{G,T}^{j,i} | d_{G,T}^{1:(j-1),i}, \mathcal{D}_f^i) \right) \quad (1)$$

where $d_{G,T}^{i,j} \in \mathcal{D}_{G,T}^i$, the superscript i is the index in a batch, the superscript j indicates the position of an element $d_{G,T}^{i,j}$ in the sequence $\mathcal{D}_{G,T}^i$, n_{batch} is the number of elements in the batch, and n_{seq}^i is the number of elements in the sequence of example i in the batch. This objective differs from the traditional LLM fine-tuning objective of maximizing the average likelihood of the next token for each token in a batch of sequences as described in equation 2;

$$Loss = \frac{1}{\sum_{i=1}^{n_{batch}} n_{seq}^i} \sum_{i=1}^{n_{batch}} \left(-\log p_{\theta}(d_{G,T}^{1,i} | \mathcal{D}_f^i) + \sum_{j=2}^{n_{seq}^i} -\log p_{\theta}(d_{G,T}^{j,i} | d_{G,T}^{1:(j-1),i}, \mathcal{D}_f^i) \right) \quad (2)$$

The difference between the two objectives is the red term in both equations. We can view the reciprocal of this term as a weight placed on the importance of an example in the training set with respect to the training objective. For the standard training objective in equation 2, the term $\sum_{i=1}^{n_{batch}} n_{seq}^i$ in red changes from batch to batch, because batches contain randomly sampled examples which differ in sequence length. This means across batches, examples in the data set are weighted differently by the loss function. For the same reason, the loss of an example in the training set will be weighted differently across epochs based on the batch it falls in. Note that this is not an issue when training or fine-tuning on text data, because in most cases all training examples in a batch have the same length. A comparison in the performance of models trained with these objectives can be seen in Tables 2 and 3. The baseline `regen` only differs from our proposed model `SGG-LLM` without message passing in that it was trained using equation 2 instead of 1. We also report additional experiments shown Appendix D which suggests what is causing the difference in performance.

¹We used GraphSAGE (15) in all experiments

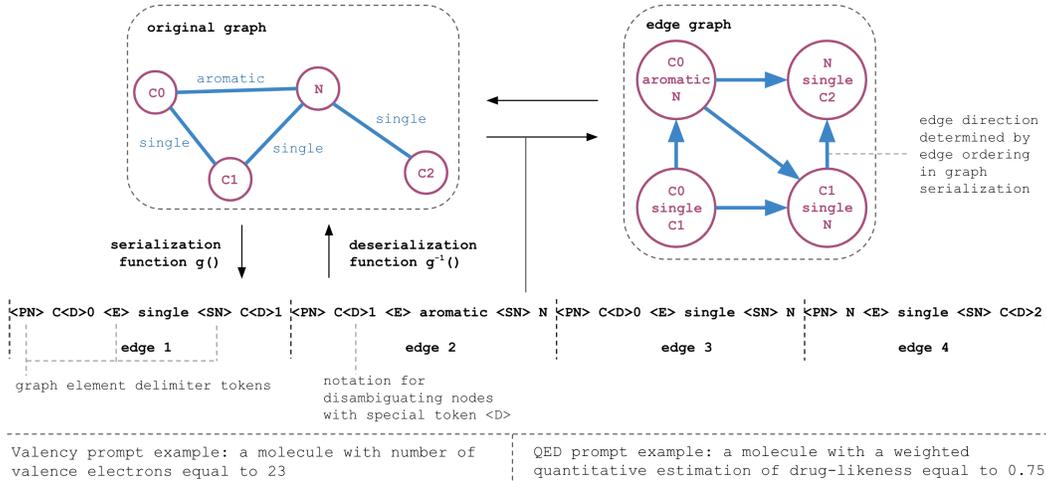


Figure 2: Top: the correspondence between a graph and its edge graph. Middle: the graph’s bag-of-edges serialisation with special tokens <PN>, <SN>, <E> and <D>. Bottom: Examples of functional requirements used for condition graph generation for the molecule above.

3.2 NODE DISAMBIGUATION DURING SERIALIZATION

One special case of graph-structured data, which is not handled by prior work on generating serialized graphs, is when a subset of nodes in a graph are described by the same feature string². However, this special case occurs in many domains, such as molecule data. For example, the molecule with multiple carbon atoms depicted in Figure 2 top row. To address this issue, we add a disambiguation token <D> to the feature strings of nodes with identical feature strings followed by a unique integer. An example of node disambiguation is shown the bottom row of Figure 2, where the three carbon atoms’ feature strings are modified from C, C, C to C<D>0, C<D>1, C<D>2. With the addition of a disambiguation token, the serialization method becomes expressive enough to reversibly serialize any graph. In conjunction with this serialization function and special token <D>, we also add special tokens <PN> (predecessor node) <E> (edge) and <SN> (successor node) to the tokenizer of the pre-trained LLM, as well as additional rows to the embedding layer of the LLM for all four special tokens before fine-tuning. See our publicly available code for implementations of the proposed serialization $g(\cdot)$ and deserialization functions $g^{-1}(\cdot)$.

3.3 MESSAGE PASSING WITHIN A LANGUAGE MODEL

Our proposed architecture requires passing information from an LLM layer to an MP layer, then to the following LLM layer, and so on as shown in block B in Figure 1. As such, we need to convert the serialized graph representation outputted by an LLM layer into a representation of the graph that can be fed into an MP layer and vice versa. The inputs to an MP layer are node feature vectors $\mathcal{D}_{V,F} \in \mathbb{R}^{N \times H}$ for each node in a graph, where H is the dimensionality of a feature vector, and the graph’s adjacency matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$. An obvious choice for constructing $\mathcal{D}_{V,F}$ would be simply using the feature vectors of $\mathcal{D}_{G,T}$. However, in $\mathcal{D}_{G,T}$, each node in the graph is described by multiple tokens and consequently multiple feature vectors in $\mathcal{D}_{G,F}$. For example, in Figure 2, the predecessor node of edge 1 is described by four tokens [$\langle \text{PN} \rangle$, C, <D>, 0]. We can calculate a single node vector from its multiple token feature vectors in $\mathcal{D}_{G,F}$ by selecting the feature vector of the last element describing the node in $\mathcal{D}_{G,F}$ because the last vector contains information about all previous tokens in the sequence including all those describing the node because of the causal attention layers in an LLM.

²Prior work on generating graphs conditioned on text focused on tasks requiring graphs that did not contain sets of nodes with the same feature string (8; 27; 33)

In addition, if a node has a degree greater than one, it will be described more than once in $\mathcal{D}_{\mathcal{G}}$ ³. While a node may occur more than once in a serialized graph, using our proposed serialization method, an edge will only occur once. In a graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, let there be a set e^k for each edge $(v_i, v_j) \in \mathcal{E}$ where the superscript k indicates the position of the edge in $\mathcal{D}_{\mathcal{G}}$ and the set contains the two nodes the edge connects $\{v_i, v_j\} \in e^k$. For the graph \mathcal{G} and graph serialization $\mathcal{D}_{\mathcal{G}}$ pair depicted in Figure 2, e^2 would be the set of nodes $\{C < D > 1, N\}$ contained in edge 2 because edge 2 is the second edge ($k = 2$) described in the serialization $\mathcal{D}_{\mathcal{G}}$. As a slight abuse of notation, we denote elements of a graph’s edge set \mathcal{E} either by (v_i, v_j) or by e^k . Hence, we define a new graph $\mathcal{G}_{edge} = \{\mathcal{V}_{edge}, \mathcal{E}_{edge}\}$ where the edges in the original graph are treated as nodes $\mathcal{V}_{edge} = \mathcal{E}$ and nodes in the original graph define edges in the new graph $\mathcal{E}_{edge} = \{(e^j, e^k) | e^j, e^k \in \mathcal{E} \wedge e^j \cap e^k \neq \emptyset \wedge j \neq k\}$. We call this new graph an edge graph. Figure 2 shows an example of transforming a graph into its edge graph where the four edges in the original graph become the nodes in the edge graph. Let $f_{index} : \mathbb{Z}^+ \rightarrow \mathbb{Z}^+$ be a mapping from the index of an edge in the sequence $\mathcal{D}_{\mathcal{V}_{edge}, F}$ to the index of the last token describing the edge in $\mathcal{D}_{\mathcal{G}, F}$. To construct the sequence of feature vectors for the nodes in an edge graph \mathcal{G}_{edge} , for each MP layer we use the last token’s feature vector of each edge in the serialized graph, so that $d_{\mathcal{V}_{edge}, F}^k = d_{\mathcal{G}, F}^{f_{index}(k)}$.

MP layers pass information between nodes in a graph by aggregating feature vectors from neighboring nodes. To ensure that an MP layer does not pass information backwards in $\mathcal{D}_{\mathcal{G}, F}$, the MP layer should only aggregate information for a node from nodes that are earlier in the sequence as depicted in causal graph attention in block D in Figure 1. causal attention, shown in block C in Figure 1, only attends to elements previous to the query element in the sequence. graph attention only attends to elements in the sequence which are neighboring nodes to the query element in the sequence based on a reference graph. causal graph attention respects both of these constraints. We propose MP layers to pass information on a graph’s edge graph. To ensure information is not passed backwards in $\mathcal{D}_{\mathcal{G}, F}$ we must add an additional constraint when constructing the edge set of an edge graph which is $\forall (e^j, e^k) \in \mathcal{E}_{edge}, k > j$. This constraint has been applied to the edge graph shown in Figure 2. After passing an edge graph through an MP layer, we add each resultant edge feature vector to the token feature vector immediately after its description in the serialized graph as described in equation 3;

$$d_{\mathcal{G}, F}^t = \begin{cases} d_{\mathcal{G}, F}^t + d_{\mathcal{V}_{edge}, F}^k \tanh a, & \text{if } f_{index}(k) = t - 1; \\ d_{\mathcal{G}, F}^t, & \text{otherwise.} \end{cases} \quad (3)$$

During the development of this method, we found it necessary to multiply element-wise the output of an MP layer by a gating term $\tanh a$ when incorporating the output back into $\mathcal{D}_{\mathcal{G}, F}$ as shown in equation 3. a is a learned gating parameter of the MP layer initialized to 0. Without the gating term, we could not fine-tune an LLM with message passing incorporated to generate valid graphs as demonstrated by the results shown in Appendix D. We hypothesize this is because the gating term allows an LLM to gradually learn to take into account graph structure. Without it, the model most likely starts fine-tuning at a position on the loss landscape from which it cannot converge to a useful local minima. The use of the gating term is inspired by the work of (1), which faced a similar issue.

4 RELATED WORK

4.1 INCORPORATING OTHER MODALITIES INTO TEXT GENERATION

This work proposes incorporating graph structure, an additional modality, into an LLM’s text generation process (where the generated text is a serialized graph). There have been many works on incorporating other modalities into the text generation process. Speech-to-text can be posed as a sequence to sequence problem (16) and so similar methods to those used for text-to-text generation have been used to incorporate the modality of sound into language generation (2; 30; 37). Another widely studied problem is image-to-text generation, where most works focus on calculating dense sequence representations of images (21; 22; 24; 29; 40; 1) to pose image-to-text generation as a

³This makes it difficult to introduce graph information into the early elements in the graph serialization sequence. For a node with a degree greater than one, if we passed the feature vector outputted from an MP layer into its earliest instance in the graph serialization, we would be passing information backwards in the sequence

sequence-to-sequence generation problem as well. From this field of work, our proposed method is inspired by (1) which incorporates image information into the text generation process by representing sequences of images as sequences of dense vectors and interleaving special layers for combining the modalities of image and text between language modelling layers in an LLM. Inspired by this approach, we propose interleaving additional layers in-between language modelling layers in a pretrained LLM and incorporating additional tokens for representing the other modality. Unlike (1), we incorporate the modality of graphs as opposed to images into the generation process, and use a modified LLM to generate graphs rather than text.

4.2 GENERATING GRAPHS CONDITIONED ON TEXT

In the field of natural language processing, there have been many works on parsing natural language into syntax trees, some of which pose the problem as a serialized graph generation conditioned on text (9; 36). However, these works use a domain specific serialization method and do not incorporate graph structure into their proposed models’ generation processes. There is one recent work (33) focused on generating explanation graphs from text which proposes a contrastive learning objective for fine-tuning the LLM T5 (31) to generate graphs, however the contrastive objective requires additional labels created by an expert. There is also a recent work published on generating protein graphs conditioned on text (25), however it cannot generalize past the domain of protein graph generation because the proposed method generates sequences of characters describing amino acid sequences.

Most prior work on text to graph generation focuses on generating knowledge graphs from text (8; 12; 14; 26; 27; 38). The best performing method on benchmark tasks is `regen` (8), which uses the LLM T5 (31) to generate knowledge graphs and is fine-tuned using the standard objective of predicting the next token in a sequence (equation 2). The second best performing method on benchmark knowledge graph generation tasks is `grapher` (27) which uses a combination of T5 and a recurrent neural network to generate graphs. This method differs from ours in the use of an additional model besides the LLM to predict edges and generate edge features as well as a training objective similar to equation 2.

5 EXPERIMENTS

Prior work on text graph generation conditioned on text does not provide an experimental design for evaluating methods that generate graphs based on functional descriptions. To do so, a dataset containing (functional requirements, graph) pairs is required for a task where it is possible to automatically calculate the functional properties of newly generated graphs. To generate such a dataset, we took the open graph benchmark large scale molecule dataset PCQM4M (20) and used the open source Python package `rdkit` (13) to generate two functional descriptions for each molecule: number of valency electrons in it and quantitative estimated-likeness (QED) (3). These are commonly used functional properties of a molecule’s graph and are described in more detail in Appendix C.

5.1 DATASETS

For evaluation we created two datasets, each dataset is composed of all the molecules in the original PCQM4M (20) dataset with less than 20 atoms, which is still more than 2 million examples. One of the two datasets contains (number of valency electrons functional description, graph) pairs and is called Valency dataset. The other dataset contains (QED functional description, graph) pairs and is called QED dataset. To evaluate the impact of amount of data on model performance, we created three subsets for each of the datasets. Each subset had 1000 randomly selected molecules in its test and validation sets. Importantly, these were the same across the three subsets and then the training sets of the three subsets were another 25,000, 100,000 and 400,000 randomly chosen examples. The training set of the smaller subsets was contained in the training set of the larger subset sets. See Figure 2 for examples of input output pairs for the constructed data sets.

	Minimum	Maximum	Mean	Standard Deviation	Median	Interquartile Range
Number of Valency Electrons	2	122	77.3	13.3	80	70 - 88
QED	0.06	0.98	0.764	0.133	0.78	0.68 - 0.88

Table 1: Summary statistics about functional requirements in datasets used in experiments. The standard deviations reported for both data sets provide context for evaluating Table 2 below.

Model	Message Passing	Training Set Size	Mean Absolute Error	
			QED	Valency
SGG-LLM	none	100,000	0.044 ± 0.011	0.060 ± 0.018
SGG-LLM	edges	100,000	0.036 ± 0.005	0.035 ± 0.014
SGG-LLM	correspondences	100,000	0.039 ± 0.007	0.045 ± 0.017
SGG-LLM	none	25,000	0.062 ± 0.008	1.703 ± 0.074
SGG-LLM	none	400,000	0.020 ± 0.001	0.076 ± 0.034
grapher	N/A	100,000	0.157 ± 0.004	1.268 ± 0.229
regen	N/A	100,000	0.149 ± 0.018	2.282 ± 1.156

Table 2: Model performance in terms of MAE on QED and Valency datasets. The **first**, **second**, and **third** best performing models are highlighted using the colors shown here. Experiments were repeated three times to estimate standard error. All fine-tuned variants of SGG-LLM outperform baselines by a statistically significant margin. Models below the boldface line are from prior work.

5.2 METRICS AND EVALUATION

To evaluate models, we calculated three metrics on the test sets of the datasets described above: parsability, diversity, and most importantly mean absolute error (MAE) with respect to the conditional functional property. A generated graph is parsable if no error is thrown when calculating its functional property. If a molecule is parsable (i.e. it follows correct serialization syntax and doesn’t violate basic laws of physics), it is given a score of 1, otherwise 0. In all experiments we use the following metrics:

- *Parsability* is the mean parsability score of samples in the test set.
- *MAE* is the mean absolute error between the generated graph’s functional property value and the requested property value averaged over samples in the test set. To interpret the magnitudes reported for MAE results see Table 1 describing some summary statistics about functional properties of graphs in the datasets used in experiments.
- *Diversity* is a measure of the multimodality of the distribution $p_{\theta}(\mathcal{D}_G|\mathcal{D}_f)$ that the LLM learns. A sampled graph is assigned a diversity score of 1, if it, a second sampled graph and the ground truth do not share the same node and edge sets; and is assigned a score of 0 otherwise. Diversity is the average of diversity scores of samples in the test set.

5.3 CURRENT STATE-OF-THE-ART

We implemented current state-of-the-art models from prior work as baselines: specifically we implemented the version of `grapher` (27), which generates edge features instead of classifying them and `regen` (8). We had to make two modifications to both approaches so they could generalize to molecule data: we added node disambiguation from Section 3.2 to their serialization methods and updated their language models to a more recent model BLOOM (34), which is the same LLM used in experiments by our proposed approach. See Appendix A for a description of the training process used to train `grapher`, `regen` and variants of our proposed approach.

6 RESULTS

On the task of generating graphs to meet functional requirements, the ideal model can generate a diverse set of parsable graphs with functional properties equal to the requested functional property. Tables 2 and 3 describe the results of our proposed approach on the QED and Valency datasets. Our

Model	Message Passing	Training Set Size	Parsability		Diversity	
			QED	Valency	QED	Valency
SGG-LLM	none	100,000	1.000 ± 0.000	0.999 ± 0.001	0.845 ± 0.017	0.506 ± 0.031
SGG-LLM	edges	100,000	0.998 ± 0.001	0.999 ± 0.000	0.836 ± 0.029	0.540 ± 0.016
SGG-LLM	correspondences	100,000	0.995 ± 0.001	0.998 ± 0.000	0.839 ± 0.008	0.608 ± 0.054
SGG-LLM	none	25,000	0.997 ± 0.002	0.991 ± 0.003	0.799 ± 0.008	0.518 ± 0.078
SGG-LLM	none	400,000	0.998 ± 0.001	1.000 ± 0.000	0.857 ± 0.006	0.542 ± 0.051
grapher	N/A	100,000	1.000 ± 0.000	1.000 ± 0.000	0.745 ± 0.038	0.410 ± 0.045
regen	N/A	100,000	0.984 ± 0.008	0.991 ± 0.007	0.854 ± 0.031	0.446 ± 0.124

Table 3: Model performance in terms of parsability and diversity on QED and Valency datasets. The **first**, **second**, and **third** best performing models are highlighted using the colors shown here. Experiments were repeated three times to estimate standard error. Models below the boldface line are from prior work. All models achieved a parsability near 1.0.

proposed approach is referred to as SGG-LLM standing for serialized graph generator large language model. All variants of our approach and the baselines grapher, and regen achieve a parsability score near or at 1. The SGG-LLM variant with correspondences message passing (described in Appendix B) is another method of incorporating the MP layers into an LLM by passing messages in \mathcal{D}_G based on node correspondences rather than edges.

Results in table 2 suggest that all variants of SGG-LLM outperform baselines by a statistically significant margin, using the unpaired t-student test and a threshold p-value of 0.05, in terms of generating examples with functional properties close those requested. In addition, results suggest that all variants of SGG-LLM outperform grapher by a statistically significant margin in terms of generating a diverse set of candidates on the Valency dataset. Finally, the two variants of SGG-LLM that use MP layers outperform the variant that does not utilise MP layers.

The high performance of all variants of our approach over regen suggest the importance of the proposed loss function and of weighting tokens equally across batches during training. The high performance of all variants of our approach over grapher could be for the same reasons as regen, as it is trained with a similar loss, or it could be that single model which jointly estimates the existence of nodes, edges, and their features is more effective at generating graphs than a dual module model.

To demonstrate the generality of our proposed approach beyond generating molecules, we also evaluate it on the benchmark knowledge graph generation dataset WebNLG+ 2020 (10) using a different language model T5 (31). Note this task is generating graphs conditioned on an imperative description of the graph, so is not directly linked to the focus of this paper. See Appendix E for a description of the experiments and discussion of the results.

7 DISCUSSION

The main limitation of our proposed approach is its reliance on fine-tuning a pre-trained autoregressive LLMs, which are computationally expensive, slow at generation time and require substantial computational resources even to fine-tune. This limitation would become even more difficult when applying this method to tasks containing larger graphs or graphs containing elements with long feature strings. Hence an interesting next step from this method would be using quantized pre-trained models and low rank adapters on LLM layers for more efficient fine tuning (7).

In terms of societal impact, our proposed approach might help speed up and improve processes such as drug discovery, software development and project planning. But at the same time requires oversight to ensure it is not used for nefarious applications like the design of chemical weapons. In addition, if this approach is applied to a task in the social sciences, analysis should be required to ensure that the biases learned by the model are understood and any unfair preferences learned for a certain demographic or group should be mitigated.

REFERENCES

- [1] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan Cabi, Tengda Han, Zhitao Gong, Sina Samangooei, Marianne Monteiro, Jacob Menick, Sebastian Borgeaud, Andy Brock, Aida Nematzadeh, Sahand Sharifzadeh, Mikolaj Binkowski, Ricardo Barreira, Oriol Vinyals, Andrew Zisserman, and Karen Simonyan. Flamingo: a visual language model for few-shot learning. *ArXiv*, abs/2204.14198, 2022.
- [2] Junyi Ao, Rui Wang, Long Zhou, Chengyi Wang, Shuo Ren, Yu Wu, Shujie Liu, Tom Ko, Qing Li, Yu Zhang, Zhihua Wei, Yao Qian, Jinyu Li, and Furu Wei. SpeechT5: Unified-modal encoder-decoder pre-training for spoken language processing. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5723–5738, Dublin, Ireland, May 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.acl-long.393. URL <https://aclanthology.org/2022.acl-long.393>.
- [3] Richard Bickerton, Gaia Paolini, Jérémy Besnard, Sorel Muresan, and Andrew Hopkins. Quantifying the chemical beauty of drugs. *Nature chemistry*, 4:90–8, 02 2012. doi: 10.1038/nchem.1243.
- [4] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners, 2020. URL <https://arxiv.org/abs/2005.14165>.
- [5] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. Evaluating large language models trained on code, 2021. URL <https://arxiv.org/abs/2107.03374>.
- [6] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pillai, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. Palm: Scaling language modeling with pathways, 2022. URL <https://arxiv.org/abs/2204.02311>.
- [7] Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. Qlora: Efficient finetuning of quantized llms. *arXiv preprint arXiv:2305.14314*, 2023.
- [8] Pierre L. Dognin, Inkit Padhi, Igor Melnyk, and Payel Das. Regen: Reinforcement learning for text and knowledge base generation using pretrained language models, 2021. URL <https://arxiv.org/abs/2108.12472>.

- [9] Chris Dyer, Adhiguna Kuncoro, Miguel Ballesteros, and Noah A. Smith. Recurrent neural network grammars, 2016.
- [10] Claire Gardent, Anastasia Shimorina, Shashi Narayan, and Laura Perez-Beltrachini. Creating training corpora for NLG micro-planners. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 179–188, Vancouver, Canada, July 2017. Association for Computational Linguistics. doi: 10.18653/v1/P17-1017. URL <https://aclanthology.org/P17-1017>.
- [11] Heming Ge, Mihir Sanjay Kale, Oshin Agarwal, Rami Al-Rfou, Siamak Shakeri, and Yunhsuan Sung. Machine translation aided bilingual data-to-text generation and semantic parsing. In *3rd Workshop on Natural Language Generation from the Semantic Web*, 2020.
- [12] Heming Ge, Mihir Sanjay Kale, Oshin Agarwal, Rami Al-Rfou, Siamak Shakeri, and Yunhsuan Sung. Machine translation aided bilingual data-to-text generation and semantic parsing. In *3rd Workshop on Natural Language Generation from the Semantic Web*, 2020.
- [13] Greg Landrum Hans De Winter Gianluca Sforna, deric4 (github handle). rdkit, 2023. URL <https://www.rdkit.org/>.
- [14] Qipeng Guo, Zhijing Jin, Xipeng Qiu, Weinan Zhang, David Wipf, and Zheng Zhang. Cyclegt: Unsupervised graph-to-text and text-to-graph generation via cycle training, 2020. URL <https://arxiv.org/abs/2006.04702>.
- [15] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30, 2017.
- [16] Awni Y. Hannun, Carl Case, Jared Casper, Bryan Catanzaro, Greg Diamos, Erich Elsen, Ryan Prenger, Sanjeev Satheesh, Shubho Sengupta, Adam Coates, and Andrew Y. Ng. Deep speech: Scaling up end-to-end speech recognition. *CoRR*, abs/1412.5567, 2014. URL <http://arxiv.org/abs/1412.5567>.
- [17] Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Tom Hennigan, Eric Noland, Katie Millican, George van den Driessche, Bogdan Damoc, Aurelia Guy, Simon Osindero, Karen Simonyan, Erich Elsen, Jack W. Rae, Oriol Vinyals, and Laurent Sifre. Training compute-optimal large language models, 2022. URL <https://arxiv.org/abs/2203.15556>.
- [18] Vahan Hovhannisyanyan, Peter Zachares, Yael Grushka-Cockayne, Alan Mosca, and Carlos Ledezma. Data-driven schedule risk forecasting for construction mega-projects. *Available at SSRN 4496119*, 2023.
- [19] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. Open graph benchmark: Datasets for machine learning on graphs. *arXiv preprint arXiv:2005.00687*, 2020.
- [20] Weihua Hu, Matthias Fey, Hongyu Ren, Maho Nakata, Yuxiao Dong, and Jure Leskovec. Ogb-lsc: A large-scale challenge for machine learning on graphs. *arXiv preprint arXiv:2103.09430*, 2021.
- [21] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc V. Le, Yunhsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. *CoRR*, abs/2102.05918, 2021. URL <https://arxiv.org/abs/2102.05918>.
- [22] Wonjae Kim, Bokyoung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without convolution or region supervision, 2021.
- [23] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL <http://arxiv.org/abs/1412.6980>.

- [24] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. Visualbert: A simple and performant baseline for vision and language. *CoRR*, abs/1908.03557, 2019. URL <http://arxiv.org/abs/1908.03557>.
- [25] Shengchao Liu, Yutao Zhu, Jiarui Lu, Zhao Xu, Weili Nie, Anthony Gitter, Chaowei Xiao, Jian Tang, Hongyu Guo, and Anima Anandkumar. A text-guided protein design framework, 2023. URL <https://arxiv.org/abs/2302.04611>.
- [26] Yaojie Lu, Qing Liu, Dai Dai, Xinyan Xiao, Hongyu Lin, Xianpei Han, Le Sun, and Hua Wu. Unified structure generation for universal information extraction, 2022. URL <https://arxiv.org/abs/2203.12277>.
- [27] Igor Melnyk, Pierre Dognin, and Payel Das. Knowledge graph generation from text, 2022. URL <https://arxiv.org/abs/2211.10511>.
- [28] Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback, 2022. URL <https://arxiv.org/abs/2203.02155>.
- [29] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. *CoRR*, abs/2103.00020, 2021. URL <https://arxiv.org/abs/2103.00020>.
- [30] Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. Robust speech recognition via large-scale weak supervision, 2022.
- [31] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(1), jan 2020. ISSN 1532-4435.
- [32] Samyam Rajbhandari, Jeff Rasley, Olatunji Ruwase, and Yuxiong He. Zero: Memory optimization towards training A trillion parameter models. *CoRR*, abs/1910.02054, 2019. URL <http://arxiv.org/abs/1910.02054>.
- [33] Swarnadeep Saha, Prateek Yadav, and Mohit Bansal. Explanation graph generation via pre-trained language models: An empirical study with contrastive learning, 2022. URL <https://arxiv.org/abs/2204.04813>.
- [34] Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, Jonathan Tow, Alexander M. Rush, Stella Biderman, Albert Webson, Pawan Sasanka Ammanamanchi, Thomas Wang, Benoît Sagot, Niklas Muennighoff, Albert Villanova del Moral, Olatunji Ruwase, Rachel Bawden, Stas Bekman, Angelina McMillan-Major, Iz Beltagy, Huu Nguyen, Lucile Saulnier, Samson Tan, Pedro Ortiz Suarez, Victor Sanh, Hugo Laurençon, Yacine Jernite, Julien Launay, Margaret Mitchell, Colin Raffel, Aaron Gokaslan, Adi Simhi, Aitor Soroa, Alham Fikri Aji, Amit Alfassy, Anna Rogers, Ariel Kreisberg Nitzav, Canwen Xu, Chenghao Mou, Chris Emezue, Christopher Klamm, Colin Leong, Daniel van Strien, David Ifeoluwa Adelani, Dragomir Radev, Eduardo González Ponferrada, Efrat Levkovizh, Ethan Kim, Eyal Bar Natan, Francesco De Toni, Gérard Dupont, Germán Kruszewski, Giada Pistilli, Hady Elsahar, Hamza Benyamina, Hieu Tran, Ian Yu, Idris Abdulmumin, Isaac Johnson, Itziar Gonzalez-Dios, Javier de la Rosa, Jenny Chim, Jesse Dodge, Jian Zhu, Jonathan Chang, Jörg Froberg, Joseph Tobing, Joydeep Bhattacharjee, Khalid Almubarak, Kimbo Chen, Kyle Lo, Leandro Von Werra, Leon Weber, Long Phan, Loubna Ben allal, Ludovic Tanguy, Manan Dey, Manuel Romero Muñoz, Maraim Masoud, María Grandury, Mario Šaško, Max Huang, Maximin Coavoux, Mayank Singh, Mike Tian-Jian Jiang, Minh Chien Vu, Mohammad A. Jauhar, Mustafa Ghaleb, Nishant Subramani, Nora Kassner, Nurulaqilla Khamis, Olivier Nguyen, Omar Espejel, Ona de Gibert, Paulo Villegas, Peter Henderson, Pierre Colombo, Priscilla Amuok, Quentin Lhoest, Rheza Harliman, Rishi Bommasani, Roberto Luis López, Rui Ribeiro, Salomey Osei, Sampo Pyysalo, Sebastian Nagel, Shamik Bose, Shamsuddeen Hassan Muhammad, Shanya Sharma, Shayne Longpre, Somaieh

Nikpoor, Stanislav Silberberg, Suhas Pai, Sydney Zink, Tiago Timponi Torrent, Timo Schick, Tristan Thrush, Valentin Danchev, Vassilina Nikoulina, Veronika Laippala, Violette Lepercq, Vrinda Prabhu, Zaid Alyafeai, Zeerak Talat, Arun Raja, Benjamin Heinzerling, Chenglei Si, Davut Emre Taşar, Elizabeth Salesky, Sabrina J. Mielke, Wilson Y. Lee, Abheesht Sharma, Andrea Santilli, Antoine Chaffin, Arnaud Stiegler, Debajyoti Datta, Eliza Szczechla, Gunjan Chhablani, Han Wang, Harshit Pandey, Hendrik Strobelt, Jason Alan Fries, Jos Rozen, Leo Gao, Lintang Sutawika, M Saiful Bari, Maged S. Al-shaibani, Matteo Manica, Nihal Nayak, Ryan Teehan, Samuel Albanie, Sheng Shen, Srulik Ben-David, Stephen H. Bach, Taewoon Kim, Tali Bers, Thibault Fevry, Trishala Neeraj, Urmish Thakker, Vikas Raunak, Xiangru Tang, Zheng-Xin Yong, Zhiqing Sun, Shaked Brody, Yallow Uri, Hadar Tojarieh, Adam Roberts, Hyung Won Chung, Jaesung Tae, Jason Phang, Ofir Press, Conglong Li, Deepak Narayanan, Hatim Bourfoune, Jared Casper, Jeff Rasley, Max Ryabinin, Mayank Mishra, Minjia Zhang, Mohammad Shoeybi, Myriam Peyrounette, Nicolas Patry, Nouamane Tazi, Omar Sanseviero, Patrick von Platen, Pierre Cornette, Pierre François Lavallée, Rémi Lacroix, Samyam Rajbhandari, Sanchit Gandhi, Shaden Smith, Stéphane Requena, Suraj Patil, Tim Dettmers, Ahmed Baruwa, Amanpreet Singh, Anastasia Cheveleva, Anne-Laure Ligozat, Arjun Subramonian, Aurélie Névél, Charles Lovering, Dan Garrette, Deepak Tunuguntla, Ehud Reiter, Ekaterina Taktasheva, Ekaterina Voloshina, Eli Bogdanov, Genta Indra Winata, Hailey Schoelkopf, Jan-Christoph Kalo, Jekaterina Novikova, Jessica Zosa Forde, Jordan Clive, Jungo Kasai, Ken Kawamura, Liam Hazan, Marine Carpuat, Miruna Clinciu, Najoung Kim, Newton Cheng, Oleg Serikov, Omer Antverg, Oskar van der Wal, Rui Zhang, Ruochen Zhang, Sebastian Gehrmann, Shachar Mirkin, Shani Pais, Tatiana Shavrina, Thomas Scialom, Tian Yun, Tomasz Limisiewicz, Verena Rieser, Vitaly Protasov, Vladislav Mikhailov, Yada Pruksachatkun, Yonatan Belinkov, Zachary Bamberger, Zdeněk Kasner, Alice Rueda, Amanda Pestana, Amir Feizpour, Ammar Khan, Amy Faranak, Ana Santos, Anthony Hevia, Antigona Unldreaj, Arash Aghagol, Arezoo Abdollahi, Aycha Tammour, Azadeh HajiHosseini, Bahareh Behroozi, Benjamin Ajibade, Bharat Saxena, Carlos Muñoz Ferrandis, Danish Contractor, David Lansky, Davis David, Douwe Kiela, Duong A. Nguyen, Edward Tan, Emi Baylor, Ezinwanne Ozoani, Fatima Mirza, Frankline Ononiwu, Habib Rezanejad, Hessian Jones, Indrani Bhattacharya, Irene Solaiman, Irina Sedenko, Isar Nejadgholi, Jesse Passmore, Josh Seltzer, Julio Bonis Sanz, Livia Dutra, Mairon Samagaio, Maraim Elbadri, Margot Mieskes, Marissa Gerchick, Martha Akinlolu, Michael McKenna, Mike Qiu, Muhammed Ghauri, Mykola Burynok, Nafis Abrar, Nazneen Rajani, Nour Elkott, Nour Fahmy, Olanrewaju Samuel, Ran An, Rasmus Kromann, Ryan Hao, Samira Alizadeh, Sarmad Shubber, Silas Wang, Sourav Roy, Sylvain Viguier, Thanh Le, Tobi Oyebade, Trieu Le, Yoyo Yang, Zach Nguyen, Abhinav Ramesh Kashyap, Alfredo Palasciano, Alison Callahan, Anima Shukla, Antonio Miranda-Escalada, Ayush Singh, Benjamin Beilharz, Bo Wang, Caio Brito, Chenxi Zhou, Chirag Jain, Chuxin Xu, Clémentine Fourier, Daniel León Perrián, Daniel Molano, Dian Yu, Enrique Manjavacas, Fabio Barth, Florian Fuhrmann, Gabriel Altay, Giyaseddin Bayrak, Gully Burns, Helena U. Vrabec, Imane Bello, Ishani Dash, Jihyun Kang, John Giorgi, Jonas Golde, Jose David Posada, Karthik Rangasai Sivaraman, Lokesh Bulchandani, Lu Liu, Luisa Shinzato, Madeleine Hahn de Bykhovetz, Maiko Takeuchi, Marc Pàmies, Maria A Castillo, Marianna Nezhurina, Mario Sànger, Matthias Samwald, Michael Cullan, Michael Weinberg, Michiel De Wolf, Mina Mihaljcic, Minna Liu, Moritz Freidank, Myungsun Kang, Natasha Seelam, Nathan Dahlberg, Nicholas Michio Broad, Nikolaus Muellner, Pascale Fung, Patrick Haller, Ramya Chandrasekhar, Renata Eisenberg, Robert Martin, Rodrigo Canalli, Rosaline Su, Ruisi Su, Samuel Cahyawijaya, Samuele Garda, Shlok S Deshmukh, Shubhanshu Mishra, Sid Kiblawi, Simon Ott, Sinee Sang-aoonsiri, Srishti Kumar, Stefan Schweter, Sushil Bharati, Tanmay Laud, Théo Gigant, Tomoya Kainuma, Wojciech Kusa, Yanis Labrak, Yash Shailesh Bajaj, Yash Venkatraman, Yifan Xu, Yingxin Xu, Yu Xu, Zhe Tan, Zhongli Xie, Zifan Ye, Mathilde Bras, Younes Belkada, and Thomas Wolf. Bloom: A 176b-parameter open-access multilingual language model, 2022. URL <https://arxiv.org/abs/2211.05100>.

- [35] Petar Veličković. Everything is connected: Graph neural networks. *Current Opinion in Structural Biology*, 79:102538, 2023.
- [36] Oriol Vinyals, Łukasz Kaiser, Terry Koo, Slav Petrov, Ilya Sutskever, and Geoffrey Hinton. Grammar as a foreign language. *Advances in neural information processing systems*, 28, 2015.
- [37] Changhan Wang, Anne Wu, Juan Miguel Pino, Alexei Baevski, Michael Auli, and Alexis Conneau. Large-scale self- and semi-supervised learning for speech translation. *CoRR*,

abs/2104.06678, 2021. URL <https://arxiv.org/abs/2104.06678>.

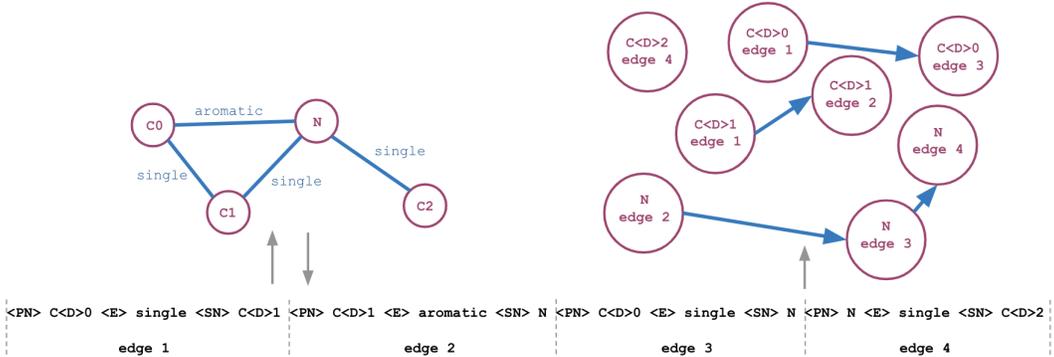
- [38] Chenguang Wang, Xiao Liu, and Dawn Song. Language models are open knowledge graphs, 2020. URL <https://arxiv.org/abs/2010.11967>.
- [39] David Weininger. Smiles, a chemical language and information system. 1. introduction to methodology and encoding rules. *J. Chem. Inf. Comput. Sci.*, 28(1):31–36, feb 1988. ISSN 0095-2338. doi: 10.1021/ci00057a005. URL <https://doi.org/10.1021/ci00057a005>.
- [40] Xiaohua Zhai, Xiao Wang, Basil Mustafa, Andreas Steiner, Daniel Keysers, Alexander Kolesnikov, and Lucas Beyer. Lit: Zero-shot transfer with locked-image text tuning. *CoRR*, abs/2111.07991, 2021. URL <https://arxiv.org/abs/2111.07991>.

A HYPERPARAMETERS AND TRAINING

For experiments on the molecule data sets, all variants of our proposed approach and the implemented baseline *grapher* (27) and *regen* (8) used the pretrained version of BLOOM (34) with 560 million parameters as their language model. For experiments on the WebNLG2020+ data set (10), all models used the pretrained version of T5 (31) with 770 million parameters as their language model. All models were trained for up to ten epochs and checkpointed based on minimizing validation loss. Model’s were trained using the ADAM optimizer (23) with a learning rate of $3e - 5$, a β_1 of 0.9, a β_2 of 0.999 and a regularization weight of $1e - 7$ as well as a linear learning rate schedule. During training, model parameters’ gradients norms were clipped to a value of 1.0. The models were trained with a batch size of 18 using stage 2 data parallelism using the method described in (32). All models were trained and evaluated on machines with 3 NVIDIA A100 GPUs. Variants of our proposed approach which incorporated message passing into the LLM by interleaving message passing layers in the LLM used a single GraphSage (15) layer in each message passing layer with an embedding size equal to the token embedding size of the LLM they were incorporated into.

B MESSAGE PASSING BETWEEN NODE CORRESPONDENCES

Most nodes appear at least twice in a *bag-of-edges* serialization. We can define a correspondence graph from a serialized graph \mathcal{D}_G by treating each instance of a node in a serialized graph \mathcal{D}_G as its own node in the correspondence graph. Then the correspondences between instances define edges in the correspondence graphs. An example of a correspondence graph is shown in the graphic below;



In the graphic, the node $C\langle D\rangle 0$ occurs more than once in \mathcal{D}_G and each instance is treated as its own node in the correspondence graph, so the node $C\langle D\rangle 0$ corresponds to $C\langle D\rangle 0$ edge 1 and to $C\langle D\rangle 0$ edge 3 in the correspondence graph. Edges are constructed by connecting nodes in the correspondence graph which correspond to the same node in the original graph i.e. in the graphic $C\langle D\rangle 0$ edge 1 is connected to $C\langle D\rangle 0$ edge 3 because they refer to the node $C\langle D\rangle 0$ in the original graph. Nodes in the correspondence graph are only connected if they are adjacent occurrences in the serialized graph. By constructing edges such that they always point from an earlier instance of a node to a later instance we ensure the message passing layer does not pass information backwards in the serialized graph sequence at training time. As additional method, we propose incorporating MP layers into an LLM where the MP layers pass information based on a graph’s correspondence graph.

C FUNCTIONAL DESCRIPTIONS OF MOLECULES

We used two functional descriptions of molecules to generate datasets for our experiments - number of valency electrons and QED. The number of valency electrons in a molecule is the sum of the number of valency electrons in its atoms. Importantly, this property is a result of only a graph’s node composition, but not its full structure. To generate descriptions of a functional property of a graph’s full structure, we calculate a metric called quantitative estimated drug-likeness (QED) (3). QED quantifies multiple properties of a molecule which are attractive for applications in drug design and then calculates a single metric from these properties using a weighted logarithmic sum. The properties used to calculate QED are molecular weight, octanol–water partition coefficient, topological polar surface area, number of hydrogen bond donors and acceptors, the number of aromatic rings and rotatable bonds, and the presence of unwanted chemical functionalities. For experiments we set the weight of the properties: molecular weight, number of hydrogen bond donors and acceptors, and presence of unwanted chemical functionalities to zero in the QED calculation, because they are mainly determined by node composition instead of the entire graph structure. The functional descriptions generated were of the format "a molecule with number of valence electrons equal to ..." and "a molecule with a weighted quantitative estimation of drug-likeness equal to ...".

D ADDITIONAL RESULTS OF MOLECULE GENERATION EXPERIMENTS

Below, we provide three tables describing the full results of experiments on the molecule data sets including ablation studies to empirically justify some design choices in our proposed method. Our proposed model is referred to as SGG-LLM. There are three additional models in the tables below; 1) SGG-LLM w/out fine-tuning, 2) SGG-LLM w/ special loss, and 3) SGG-LLM with edge based message passing without a gating term. The table below describing performance of models in terms of the parsability of generated examples shows that SGG-LLM w/out fine-tuning and SGG-LLM with edge based message passing without a gating term both did not produce a single parsable example on molecule datasets’ test sets. These results suggests that both fine-tuning and a gating term (when incorporating message passing into an LLM) are required to achieve good performance with our proposed method. Note if a model cannot generate parsable examples, then the metric mean absolute error cannot be calculated and the diversity of generated examples is not a useful measure of model performance. Consequently, we do not report mean absolute error or diversity for SGG-LLM w/out fine-tuning and SGG-LLM with edge based message passing without a gating term.

Model	Message Passing	Training Set Size	Parsability	
			QED	Valency
SGG-LLM	none	100,000	1.000 ± 0.000	0.999 ± 0.001
SGG-LLM w/out fine-tuning	none	100,000	0.000 ± 0.000	0.000 ± 0.000
SGG-LLM w/ special loss	none	100,000	0.986 ± 0.003	-
SGG-LLM	edges	100,000	0.998 ± 0.001	0.999 ± 0.000
SGG-LLM w/out gating term	edges	100,000	0.000 ± 0.000	0.000 ± 0.000
SGG-LLM	correspondences	100,000	0.995 ± 0.001	0.998 ± 0.000
SGG-LLM	none	25,000	0.997 ± 0.002	0.991 ± 0.003
SGG-LLM	none	400,000	0.998 ± 0.001	1.000 ± 0.000
grapher	N/A	100,000	1.000 ± 0.000	1.000 ± 0.000
regen	N/A	100,000	0.984 ± 0.008	0.991 ± 0.007

The model SGG-LLM w/ special loss is a version of SGG-LLM without message passing that was trained using equation 1, but instead of letting $n_{batch} = \text{batch size}$ as proposed in section 3.1, we equal $n_{batch} = E[\sum_{i=1}^{n_{batch}} n_{seq}^i]$ which is the expected value of the differing term in equation 2 from equation 1. This is to determine whether it is the magnitude of the differing term causing the difference in performance or the fact that $\sum_{i=1}^{n_{batch}} n_{seq}^i$ changes from batch to batch. In experiments, some summary statistics for the term $\sum_{i=1}^{n_{batch}} n_{seq}^i$ in equation 2 were mean = 4053, max = 4676, median = 4060, and inter-quartile range = 3915 – 4199. There was a marginal difference in performance in terms of all three metrics when comparing the performance of SGG-LLM w/ special

loss and SGG-LLM without message passing. The marginal difference in performance suggests that the changing weighting between batches and across epochs is what hurts the performance of models trained with the objective described in equation 1. *regen* is a model trained with equation 1.

Model	Message Passing	Training Set Size	Mean Absolute Error	
			QED	Valency
SGG-LLM	none	100,000	0.044 ± 0.011	0.060 ± 0.018
SGG-LLM w/ special loss	none	100,000	0.046 ± 0.004	-
SGG-LLM	edges	100,000	0.036 ± 0.005	0.035 ± 0.014
SGG-LLM	correspondences	100,000	0.039 ± 0.007	0.045 ± 0.017
SGG-LLM	none	25,000	0.062 ± 0.008	1.703 ± 0.074
SGG-LLM	none	400,000	0.020 ± 0.001	0.076 ± 0.034
grapher	N/A	100,000	0.157 ± 0.004	1.268 ± 0.229
regen	N/A	100,000	0.149 ± 0.018	2.282 ± 1.156

Model	Message Passing	Training Set Size	Diversity	
			QED	Valency
SGG-LLM	none	100,000	0.845 ± 0.017	0.506 ± 0.031
SGG-LLM w/ special loss	none	100,000	0.831 ± 0.003	-
SGG-LLM	edges	100,000	0.836 ± 0.029	0.540 ± 0.016
SGG-LLM	correspondences	100,000	0.839 ± 0.008	0.608 ± 0.054
SGG-LLM	none	25,000	0.799 ± 0.008	0.518 ± 0.078
SGG-LLM	none	400,000	0.857 ± 0.006	0.542 ± 0.051
grapher	N/A	100,000	0.745 ± 0.038	0.410 ± 0.045
regen	N/A	100,000	0.854 ± 0.031	0.446 ± 0.124

E RESULTS OF KNOWLEDGE GRAPH GENERATION EXPERIMENTS

To demonstrate the generality of our proposed approach beyond generating molecules, we also evaluate it on the benchmark knowledge graph generation dataset WebNLG+ 2020 (10) using a different language model T5 (31). Like in prior work, model performance is evaluated using the metrics of F1-score, precision and recall when comparing a generated graph to the ground truth knowledge graph. See (8) for a more detailed explanation of these metrics. We compare our proposed method to the three best performing models on this data set; *regen* (8), *grapher* (27), and *bt5* (11).

On the WebNLG+ 2020 data set, results in the table below suggest that the incorporating message passing into an LLM is not useful to the task of knowledge graph generation from an imperative description, but also that using edge message passing does not degrade performance. Interestingly, SGG-LLM without message passing was able to achieve state-of-the-art performance on the benchmark task of generating triples in a knowledge graph. *regen*'s implementation for experiments on WebNLG2020+ was identical to SGG-LLM without message passing, except that SGG-LLM was trained with a different training objective (see equation 1 in the main paper), a more aggressive learning rate and a linear learning rate schedule. So the state-of-the-art performance of SG-LLM on WebNLG2020+ may be attributed to better hyperparameter selection or the modified training objective.

Model	Exact			Partial			Strict		
	F1	Prec.	Rec.	F1	Prec.	Rec.	F1	Prec.	Rec.
bt5	0.682	0.670	0.701	0.713	0.700	0.736	0.675	0.663	0.695
regen	0.723	0.714	0.738	0.767	0.755	0.788	0.720	0.713	0.735
grapher	0.683	0.675	0.695	0.713	0.702	0.730	0.681	0.673	0.693
SGG-LLM none	0.747 ± 0.005	0.743 ± 0.003	0.765 ± 0.007	0.779 ± 0.006	0.771 ± 0.005	0.798 ± 0.004	0.726 ± 0.005	0.720 ± 0.007	0.731 ± 0.012
SGG-LLM edges	0.754 ± 0.003	0.748 ± 0.002	0.762 ± 0.004	0.786 ± 0.001	0.779 ± 0.000	0.795 ± 0.003	0.733 ± 0.005	0.729 ± 0.004	0.741 ± 0.003
SGG-LLM correspo- ndences	0.743 ± 0.008	0.738 ± 0.009	0.752 ± 0.005	0.774 ± 0.005	0.769 ± 0.009	0.787 ± 0.002	0.683 ± 0.067	0.693 ± 0.081	0.691 ± 0.069

Model performance on knowledge graph data set WebNLG+ 2020. Note: baseline model results do not report standard deviations because they were not reported in prior work and we felt it was more appropriate to report baseline results based their published results as opposed to reimplementing the baselines ourselves. Experiments with variants of our proposed approach were repeated three times to estimate standard error.