
Group SELFIES: A Robust Fragment-Based Molecular String Representation

Austin H. Cheng^{*†}
University of Toronto

Andy Cai^{*}
University of Toronto

Santiago Miret
Intel Labs

Gustavo Malkomes
SigOpt

Mariano Phielipp
Intel Labs

Alán Aspuru-Guzik
University of Toronto

Abstract

We introduce Group SELFIES, a molecular string representation that leverages group tokens to represent functional groups or entire substructures while maintaining chemical robustness guarantees. Molecular string representations, such as SMILES and SELFIES, serve as the basis for molecular generation and optimization in chemical language models, deep generative models, and evolutionary methods. While SMILES and SELFIES leverage atomic representations, Group SELFIES builds on top of the chemical robustness guarantees of SELFIES by enabling group tokens, thereby creating additional flexibility to the representation. Moreover, the group tokens in Group SELFIES can take advantage of inductive biases of molecular fragments that capture meaningful chemical motifs. The advantages of capturing chemical motifs and flexibility are demonstrated in our experiments, which show that Group SELFIES improves distribution learning of common molecular datasets. Further experiments also show that random sampling of Group SELFIES strings improves the quality of generated molecules compared to regular SELFIES strings. Our open-source implementation of Group SELFIES is available at <https://github.com/aspuru-guzik-group/group-selfies>, which we hope will aid future research in molecular generation and optimization.

1 Introduction

The discovery of functional molecules for drugs and energy materials is crucial to tackling global challenges in public health and climate change. Different types of generative models can suggest potential molecules to synthesize and test, but the performance of the models and molecules heavily relies on the underlying molecular representation. Several models generate molecules represented as SMILES strings [1–5], but their generated output can be invalid due to syntax errors or incorrect valency. SELFIES [6] is a molecular string representation that overcomes chemical invalidity challenges by ensuring that any string of SELFIES characters can be decoded to a molecule with valid valency. This not only makes it a natural representation for chemical language models that output molecular strings, but also for genetic algorithms such as GA+D, STONED, and JANUS [7–9] for molecular optimization.

SELFIES improves string-based molecular generation by encoding prior knowledge of valency constraints into the representation independently of the optimization method. The representation has been shown to improve distribution learning by language models [10], as well as image2string and

^{*}Equal contribution.

[†]Correspondence to: austin@cs.toronto.edu

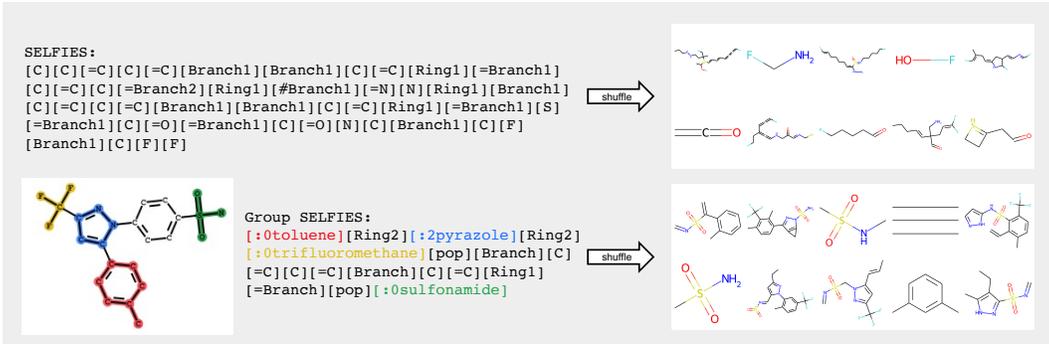


Figure 1: Visual overview of SELFIES and Group SELFIES. SELFIES is robust, so shuffling tokens around will yield new molecules with correct valency. Group SELFIES maintains robustness while adding group tokens, highlighted in color. When Group SELFIES tokens are shuffled, structures like benzene rings are more often preserved, while shuffled SELFIES strings rarely ever preserve structures. Incidentally, Group SELFIES also improves the readability of molecular string representations since chemists can see what substructures are present.

string2string translation [11, 12] and molecular generation in data-sparse regimes [13]. Additionally, simple add/replace/delete edits to SELFIES strings can generate new but similar molecules, enabling genetic algorithms that directly manipulate strings to generate molecules [7]. Alternatively, guiding these simple string edits with Tanimoto similarity can interpolate between molecules as performed in STONED by Nigam et al. [8], which can then be applied as crossover operations in genetic algorithms such as JANUS [9]. Molecular interpolation has also been used to find counterfactual decision boundaries that explain a molecular classifier’s decisions [14].

While SMILES and SELFIES represent molecules at the individual atom and bond level, human chemists typically think about molecules in terms of the substructures that they contain. Human chemists can distinguish molecular substructures based solely on the image of a molecule and induce the molecular properties those substructures usually imply. Many fragment-based generative models take advantage of this inductive bias [15–24] by constructing custom representations amenable to fragment-based molecular design. However, these approaches are not string-based, thereby losing desirable properties of string representations: easy manipulation, and direct input into established language models.

Similar to how SELFIES incorporates prior knowledge of valency constraints, we can also incorporate prior knowledge in the form of functional groups and molecular substructures into the representation. In this work, we combine the flexibility of string representations with the chemical robustness of SELFIES and the interpretability and inductive bias of fragment-based approaches into a novel string representation: Group SELFIES, a robust string representation that extends SELFIES to include tokens which represent functional groups or entire substructures.

In Section 2, we discuss how Group SELFIES fits into related research and then formally introduce the representation in Section 3. Specifically, we outline how molecules are encoded into and decoded from Group SELFIES, and we show that arbitrary Group SELFIES strings can be decoded to molecules with valid valency. The representation enables users to easily specify their own groups or extract fragments from a dataset, leveraging the wide area of cheminformatics research available there. In Section 4, we find that Group SELFIES is more compact than SMILES or SELFIES and improves distribution learning. Additionally, we compare molecules generated via randomly sampling SELFIES and Group SELFIES strings and find that Group SELFIES improves the quality of generated molecules. Molecular generation via random sampling provides greater emphasis on the representation itself by abstracting away the complexities of the type of generative method used, which we leave to future work as described in Section 5.

Overall, Group SELFIES provides the flexibility of *group representation*, the ability to represent *extended chirality* via chiral group tokens and *chemical robustness* as summarized in Table 1.

Representation	Robustness	Substructure Control	Extended Chirality	Distribution Learning
SMILES	no	no	no	~
SELFIES	yes	no	no	~
Group SELFIES	yes	yes	yes	improved

Table 1: Comparison of the capabilities of SMILES, SELFIES, and Group SELFIES. Group SELFIES provides *group representation*, representation of *extended chirality* and *chemical robustness*. Additionally as shown in Section 4, Group SELFIES improves distribution learning compared to other representations.

2 Related Work

Fragment-Based String Representations: Group SELFIES is not the first fragment-based string representation that has been proposed. Historical string representations, such as Wiswesser Line Notation (WLN) [25–27], Hayward Notation [28], and Skolnik Notation [29], all predate SMILES and represent molecules non-atomically. They use tokens that represent functional groups, such as carboxyls or phenyls, as well as ring systems. WLN strings are usually shorter and sometimes easier for trained humans to understand than SMILES, as it is easier to recognize functional groups encoded as single characters than functional groups encoded atomically. SYBYL Line Notation (SLN) [30] allows for “macro atoms” which specify multiple atoms in a substructure. The Hierarchical Editing Language for Macromolecules (HELM) [31] represents complex biomolecules by declaring monomers and then connecting them in a polymer line notation. Human-Readable SMILES [32] applies common abbreviations for chemical substituents to process and compress SMILES strings into a more human-readable format. SMILES Pair Encoding [33] breaks down SMILES strings by tokenizing them in a data-driven way that recognizes common substructures.

Learned Grammars: Data-Efficient Graph Grammar Learning (DEG) [22] is a recent approach for extracting useful formal graph grammars from small datasets of molecules. In this context, a “useful grammar” means that molecules generated by applying random applicable production rules usually have high scores. The learned production rules of the graph grammar can be thought of as similar to functional groups applied in Group SELFIES. Group SELFIES allows for flexibility and fine control of substructures, which can be extracted from any procedure including DEG.

3 Representation

3.1 SELFIES Framework

Before introducing Group SELFIES in greater detail, we summarize the primary features of SELFIES and the reasons underlying its chemical robustness. SELFIES is equipped with an encoder and a decoder. The encoder takes in a molecule and converts it to a SELFIES string, and the decoder takes in a SELFIES string and converts it to a molecule. To encode a molecule in SELFIES, one traverses its molecular graph and outputs the processed traversal as a string of SELFIES tokens. To decode a SELFIES string, one reads through the string token-by-token, building the molecular graph along the way until arriving at the finished graph. Since the encoding and decoding process alone does not guarantee chemical robustness, the SELFIES decoder further includes two important features:

1. Each token in SELFIES is overloaded to ensure that it can be interpreted sensibly in all contexts. For instance, all tokens in SELFIES can also be interpreted as numbers, which is useful when expressing branch and ring lengths.
2. SELFIES keeps track of the available valency at each step in the decoding process; if a bond would be formed that would exceed this valency, it changes the bond order or ignores the bond. For instance, when decoding [C][O][=C], adding [=C] would exceed the valency of [O], so SELFIES changes the bond order and adds [C] instead.

3.3 Groups

The primary difference between SELFIES and Group SELFIES is the addition of groups. Each group is defined as a set of atoms and bonds representing the molecular group with its *attachment points*, indicating how the group can participate in bonding. Each attachment point has a specified maximum valency, which allows us to continue tracking available valency while decoding. These attachment points are labeled by *attachment indices*, and the encoder and decoder will navigate around these attachment indices as described in Section 3.4.

Users must specify the groups they want to use using a dictionary that maps group names to groups. This tells the encoder what groups to recognize, and tells the decoder how to map group tokens to groups. We call this dictionary a “group set”, and every group set defines its own distinct instance of Group SELFIES. In particular, the decoder will not recognize a Group SELFIES string that contains group tokens not present in the current group set.

To distinguish group tokens from other tokens, we include a `:` character at the front of the token (e.g. `[:4benzene]`). All group tokens are of the form `[:S<group-name>]`, where S is the starting attachment index of the group, and `<group-name>` is any alphanumeric string that does not start with a number.

Optionally, a priority value can be specified for each group, indicating the priority with which the group should be recognized when encoding into Group SELFIES. Priority affects the Group SELFIES encoder as described in Section 3.4. For each group, one can also specify its index overload value, which is the value the group token takes when the decoder must interpret the token as a number.

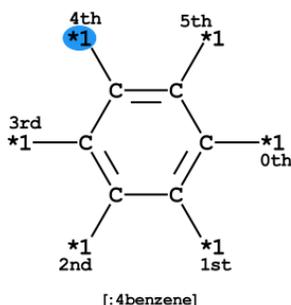


Figure 3: Representation of a possible “benzene” group and its use in a corresponding group token. The `*N` notation represents an attachment point with valency `N`. Each attachment point is labeled with its attachment index `0th`, `1st`, . . . The starting attachment point represented in the token is also highlighted.

3.4 Encoding and Decoding

Encoding: To encode a molecule in Group SELFIES, the encoder first recognizes and replaces substructure matches of groups from the molecule. By default, the encoder iterates through the group set and recognizes the largest groups first, but users can override this by specifying a priority for each group. Setting a high priority value for a group indicates that it will be recognized first when encoding into Group SELFIES, ensuring that other group replacements will not overlap with this group. This encoding strategy iterates implies that increasing the size of the group set will increase the running time of the encoder linearly. We then traverse the graph similar to the encoding process for SMILES and SELFIES, while also placing the correct tokens for tracking the attachment indices of where the encoder entered and exited a group.

Decoding: When decoding Group SELFIES, the process is essentially the same as regular SELFIES except when reading group tokens. When a group token is read by the decoder, the group set dictionary determines the corresponding group. Subsequently, all atoms of the group are placed and the main chain is connected to the starting attachment point. The decoder selects the next attachment point to branch off from by reading in the next token as a *relative index*. By adding the current attachment index to a relative index modulo the total number of attachment points in the group, the decoder selects the next attachment point. From the specified attachment point, the decoder implicitly

branches off of the group and continues traversing until a [pop] token is read. Once the branch is “popped”, the decoder returns to the group and can navigate to the next attachment point using another relative index. If the selected attachment point is occupied, then the next available attachment point is used. If all attachment points have been used up, then the group itself is immediately “popped”, returning to the most recent branchpoint before the group was placed.

We verified the robustness of Group SELFIES by encoding and decoding 25M molecules from the eMolecules database [34]. We provide a detailed example of encoding/decoding the molecule celecoxib in Appendix A.1.

Group SELFIES manages chirality differently than SMILES and SELFIES. Rather than use @-notation to specify tetrahedral chirality, all chiral centers must be specified as groups. We provide an “essential set” of 23 groups which encode all relevant chiral centers in the eMolecules database. Equipped with this essential set, every molecule can be encoded-decoded while maintaining chirality. It is also an option to not use the essential set, or only use a subset of it, depending on what chiral centers are relevant to the problem at hand. If a molecule has a chiral center not specified in the group set, then encode-decode will not preserve chirality.

3.5 Determining Fragments

Group SELFIES has a built-in flexibility for assigning the set of fragments that make up a group set. Hence, the construction of a useful group set often remains an open design choice. Users can specify groups using a SMILES-like syntax, which could be useful if one knows what groups are synthetically available or are expected to be useful for their particular design task. Fragments can also be obtained from several fragment libraries found in the literature [35–37]. Generally, a useful set of groups will appear in many molecules in the dataset and replace many atoms, with similar fragments merged together to reduce redundancy.

In our experiments, we also tested various fragmentation algorithms that extract fragments from a dataset, including a naïve technique that cleaves side chains from rings and a method based on matched molecular pair analysis [38]. Several other fragmentation algorithms from cheminformatics can be readily applied [22, 39–43] and the essential features of Group SELFIES outlined in Section 1 remain robust to different fragment discovery methods.

4 Experiments

The experiments in the subsequent sections outline some of the advantages of Group SELFIES compared to SMILES and regular SELFIES representations. Concretely, we show that: (1) Group SELFIES is shorter and more compressible than SMILES and SELFIES; (2) Group SELFIES preserves useful properties during generation; (3) Group SELFIES improves distribution learning.

4.1 Compactness

Group SELFIES strings are typically shorter than their SMILES and SELFIES equivalents when using a generic set of groups. In Figure 4, this generic set was generated by taking a random selection of 10,000 molecules from ZINC-250k [44] and fragmenting them into 30 useful groups using various algorithms (see Determining Fragments). We then combined these 30 groups with the 23 groups of the essential set. Figure 4 shows histograms of the lengths of SMILES/SELFIES/Group SELFIES strings of the entire ZINC-250k dataset. Length is the number of characters in SMILES strings, and the number of tokens in (Group) SELFIES strings. Group SELFIES strings are usually shorter than their SELFIES and SMILES counterparts because group tokens can represent multiple atoms in a molecule.

Since Group SELFIES has a larger alphabet than SMILES or SELFIES, we estimate the complexity of each representation with the compressed filesize of ZINC-250k. We find that out of all representations, Group SELFIES can be compressed the most (see Appendix A.2).

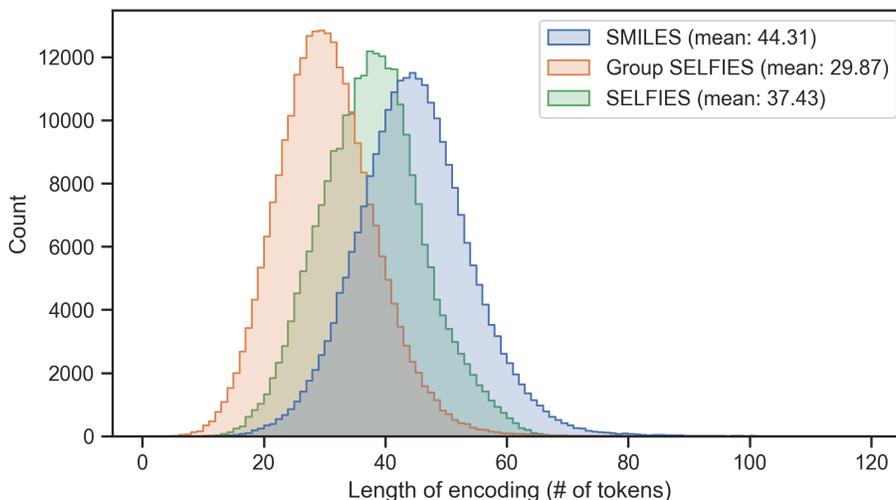


Figure 4: Histogram of lengths of SMILES, SELFIES, and Group SELFIES strings of the ZINC-250k dataset. Here, Group SELFIES uses a group set of 53 groups.

4.2 Molecular Generation

To specifically compare the suitability of the SELFIES and Group SELFIES representations for molecular generation, we use a primitive generative model which samples random strings. First, we convert a subset of $N = 100,000$ molecules from ZINC-250k into (Group) SELFIES strings. Then, we tokenize all strings and combine them into a single bag of tokens. To generate a new string, we first pick a random (Group) SELFIES string from our chosen subset and take its length l . We then randomly sample l tokens from the bag, and concatenate into a generated string. We generate N random strings for each representation. For Group SELFIES, we use the same 53 groups used for the length histogram in Section 4.1.

We show histograms of the SAScore [45] and QED [46] of molecules generated from ZINC in Figure 5. The distributions of generated Group SELFIES more closely overlap with the original ZINC dataset than the generated SELFIES, showing that even with an extremely simplistic generative model, Group SELFIES can preserve important structural information. We perform a similar analysis for a dataset of nonfullerene acceptors (NFA) [47] in Appendix A.3 and find that Group SELFIES preserves many aromatic rings in contrast to SELFIES, which rarely ever preserves aromatic rings.

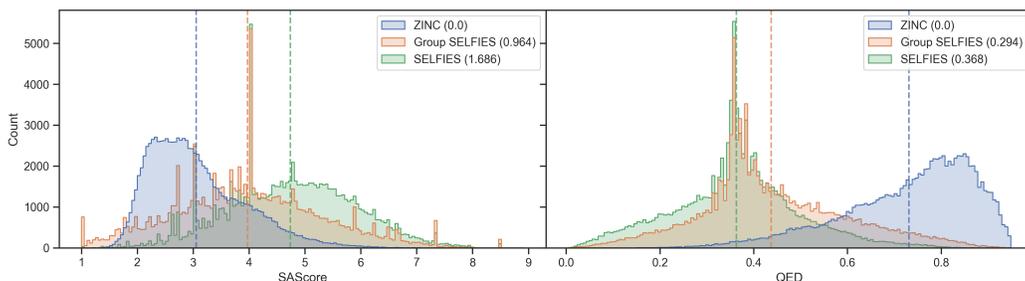


Figure 5: Molecules generated by our primitive generative model are binned by SAScore and QED. For both properties, generated Group SELFIES have greater overlap with the original ZINC distribution. Bracketed values indicate the Wasserstein distance (a measure of overlap) to the ZINC distribution. Dashed lines indicate the means.

4.3 Distribution Learning

To further quantify the effectiveness of Group SELFIES in generative models, we use the MOSES benchmarking framework [48] to evaluate variational autoencoders (VAEs) trained with both Group SELFIES and SELFIES strings. Models were trained for 125 epochs. The group set for the Group SELFIES VAE was created by fragmenting the training set provided by MOSES and selecting the 300 most diverse groups. A set of 100,000 molecules was then generated from each model and evaluated on the metrics provided by MOSES.

Model	Valid (\uparrow)	Unique@1k (\uparrow)	Unique@10k (\uparrow)	FCD (\downarrow)		SNN (\uparrow)	
				Test	TestSF	Test	TestSF
<i>Train</i>	1.0	1.0	1.0	0.008	0.4755	0.6419	0.5859
Group-VAE-125	1.0(0)	1.0(0)	0.9985(4)	0.1787(29)	0.734(109)	0.6051(4)	0.5599(3)
SELFIES-VAE-125	1.0(0)	0.9996(5)	0.9986(4)	0.6351(43)	1.3136(128)	0.6014(3)	0.5566(2)

Model	Frag (\uparrow)		Scaf (\uparrow)		IntDiv (\uparrow)	IntDiv2 (\uparrow)	Filters (\uparrow)	Novelty (\uparrow)
	Test	TestSF	Test	TestSF				
<i>Train</i>	1.0	0.9986	0.9907	0.0	0.8567	0.8508	1.0	1.0
Group-VAE-125	0.9995(0)	0.9977(1)	0.9649(21)	0.0608(65)	0.8587(1)	0.8528(1)	0.9623(7)	0.7187(11)
SELFIES-VAE-125	0.9989(0)	0.9965(1)	0.9588(15)	0.0675(37)	0.8579(1)	0.8519(1)	0.96(4)	0.7345(16)

Table 2: Group SELFIES VAE and SELFIES VAE evaluated on MOSES metrics. The Group SELFIES VAE mostly matches or outperforms the SELFIES VAE.

For most metrics, Group SELFIES performs approximately the same as SELFIES. Validity is the percentage of generated molecules that are accepted by RDKit’s parser. Uniqueness is the percentage of generated molecules that are not identical to any other generated molecule. Similarity to nearest neighbor (SNN) is the average Tanimoto similarity between generated molecules and the nearest neighbor in the reference set. Fragment similarity (Frag) is a cosine similarity based on the distribution of BRICS fragments [43] of generated and reference molecules. Scaffold similarity (Scaf) is a cosine similarity based on the distribution of Bemis-Murcko scaffolds [49] of generated and reference molecules. Internal diversity (IntDiv) measures the chemical diversity of the generated molecules using Tanimoto similarity. Filters is the fraction of generated molecules that pass filters for unwanted fragments. Novelty is the fraction of generated molecules not in the training set.

The Group SELFIES model performs especially well on the Fréchet ChemNet Distance (FCD) metric [50] when compared to SELFIES. FCD measures the difference between the activations of the penultimate layer of ChemNet (a model trained to predict the bioactivity of molecules) in the validation set and in the generated set. Due to how ChemNet was trained, the activations are likely to encode a mixture of biological and chemical properties important to drug likelihood. This makes comparing these activations more informative than comparing standard properties like logP or molecular weight, where the correlation to bioactivity is weaker and less deliberate. To visualize FCD, some indices of the penultimate activations of ChemNet are graphed in Figure 6. Generated Group SELFIES match these distributions more closely than generated SELFIES.

5 Discussion

Our experiments show that Group SELFIES has noticeable advantages compared to SMILES and SELFIES representations, including greater readability provided by the group tokens. With regards to SMILES, the primary advantage is chemical robustness. The comparison with SELFIES is more nuanced, as discussed in the section below.

5.1 Group SELFIES vs SELFIES

Substructure Control: Group SELFIES provides more fine-grained control of substructures, which creates the following advantages: (1) An important scaffold can be preserved during optimization; (2) Chiral and charged groups can be preserved during optimization, ensuring that charged tokens do not proliferate and create radicals; (3) Synthetically accessible building blocks can be chosen as groups to improve synthesizability.

Substructure Control with SELFIES: Various techniques applied to SELFIES can mitigate the challenges of preserving structure. One such example is to simply combine substrings of SELFIES

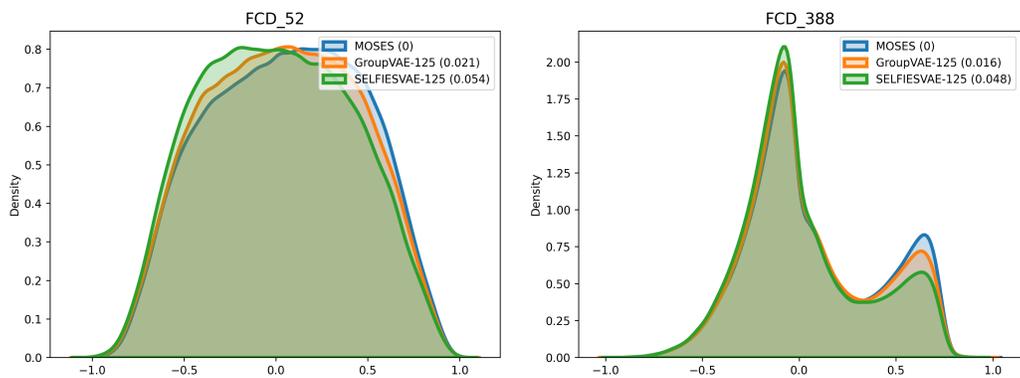


Figure 6: Distribution of some values from the second-to-last layer of ChemNet for molecules generated by Group SELFIES and SELFIES compared to the validation set. The difference in distributions is used to calculate FCD. Bracketed values in the legend represent the Wasserstein distance to the original MOSES distribution.

strings together. Indeed, further experiments in Appendix A.4 show that simply replacing all group tokens by their SELFIES substrings shows similar performance to Group SELFIES. Within the SELFIES framework, however, the first token of the inserted substring may need to be interpreted as a number, which can have cascading effects for the rest of the substring. It is also likely that upon further insertions, that substructure will not be preserved. Additionally, it is also not clear how an insertion based approach can create groups with 3 or more branches, since creating a third branch requires insertion in the middle of the group substring.

Extended Chirality: Group SELFIES is theoretically capable of representing molecules which are traditionally not able to be represented with SMILES or SELFIES, such as complexes like ferrocene and molecules with axial/helical/planar chirality. This can be done by taking the entire complex or chiral substructure and abstracting it into a group, leaving attachment points on the outside for varying functionalization. We leave this for future work.

Computational Speed: One tradeoff of Group SELFIES is that encoding and decoding is usually slower than with SELFIES, likely due to overhead of RDKit operations. The encoder is particularly slow as it relies on performing a substructure match for every group in the group set. The decoder is faster than the encoder, though still slower than the SELFIES decoder. See Appendix A.6 for timing. To improve computational performance in future work, one could exploit substructure control of Group SELFIES to reduce the number of encode/decode calls needed to obtain high performers. Additionally, the speed of encoding and decoding operations can be improved with distributed computing, since Group SELFIES is trivially parallelizable for a fixed group set.

5.2 Future Work

One promising extension of Group SELFIES is to incorporate more flexibility into the representation. In such a case, a group token can represent an entire scaffold, except without the atom identities. Other tokens can then identify the atom types on the scaffolds. This would allow optimization of the atom types while maintaining the topological structure of the scaffold. Another current limitation of Group SELFIES is that groups cannot overlap; more work is needed to develop a representation that acknowledges how groups might overlap, particularly for generating polycyclic compounds. A sequence-based representation of cellular complexes [51] or hypergraphs might be promising.

Group SELFIES can serve as a basis for a variety of generative methods that can be applied to AI-Guided materials design, including but not limited to deep generative models (VAE, GAN), evolutionary search, and reinforcement learning, especially those that apply fragment-based molecular design. In this paper we focused on the introduction of the representation and its capabilities and leave the application of generative methods to future work due to the vast variety of methods the representation could be applied to. For illustrative purposes, we already included a simple VAE benchmark in Section 4.3 and works such as JANUS [9] could be extended to leverage Group SELFIES.

Acknowledgments and Disclosure of Funding

We thank Luca Thiede, Akshat Nigam, Robert Pollice, Kjell Jorner, Gary Tom, Nathanael Kusanda, Edwin Yu, Naruki Yoshikawa, and Felix Strieth-Kalthoff for useful discussions.

Calculations testing the robustness of Group SELFIES were performed on the Niagara supercomputer at the SciNet HPC Consortium. SciNet is funded by: the Canada Foundation for Innovation; the Government of Ontario; Ontario Research Fund - Research Excellence; and the University of Toronto.

This research was partially supported by funds from Intel Corporation.

A.A.-G. thanks Anders G. Frøseth for his generous support. A.A.-G. also acknowledges the generous support of Natural Resources Canada and the Canada 150 Research Chairs program.

References

- [1] Seyone Chithrananda, Gabriel Grand, and Bharath Ramsundar. Chemberta: large-scale self-supervised pretraining for molecular property prediction. *arXiv preprint arXiv:2010.09885*, 2020.
- [2] Rafael Gómez-Bombarelli, Jennifer N Wei, David Duvenaud, José Miguel Hernández-Lobato, Benjamín Sánchez-Lengeling, Dennis Sheberla, Jorge Aguilera-Iparraguirre, Timothy D Hirzel, Ryan P Adams, and Alán Aspuru-Guzik. Automatic chemical design using a data-driven continuous representation of molecules. *ACS central science*, 4(2):268–276, 2018.
- [3] Thomas Blaschke, Josep Arús-Pous, Hongming Chen, Christian Margreitter, Christian Tyrchan, Ola Engkvist, Kostas Papadopoulos, and Atanas Patronov. Reinvent 2.0: an ai tool for de novo drug design. *Journal of chemical information and modeling*, 60(12):5918–5922, 2020.
- [4] Henry Moss, David Leslie, Daniel Beck, Javier Gonzalez, and Paul Rayson. Boss: Bayesian optimization over string spaces. *Advances in neural information processing systems*, 33: 15476–15486, 2020.
- [5] Benjamin Sanchez-Lengeling, Carlos Outeiral, Gabriel L Guimaraes, and Alan Aspuru-Guzik. Optimizing distributions over molecular space. an objective-reinforced generative adversarial network for inverse-design chemistry (organic). 2017.
- [6] Mario Krenn, Florian Häse, AkshatKumar Nigam, Pascal Friederich, and Alan Aspuru-Guzik. Self-referencing embedded strings (selfies): A 100% robust molecular string representation. *Machine Learning: Science and Technology*, 1(4):045024, 2020.
- [7] AkshatKumar Nigam, Pascal Friederich, Mario Krenn, and Alán Aspuru-Guzik. Augmenting genetic algorithms with deep neural networks for exploring the chemical space. *arXiv preprint arXiv:1909.11655*, 2019.
- [8] AkshatKumar Nigam, Robert Pollice, Mario Krenn, Gabriel dos Passos Gomes, and Alan Aspuru-Guzik. Beyond generative models: superfast traversal, optimization, novelty, exploration and discovery (stoned) algorithm for molecules using selfies. *Chemical science*, 12(20):7079–7090, 2021.
- [9] AkshatKumar Nigam, Robert Pollice, and Alán Aspuru-Guzik. Parallel tempered genetic algorithm guided by deep neural networks for inverse molecular design. *Digital Discovery*, 2022.
- [10] Daniel Flam-Shepherd, Kevin Zhu, and Alán Aspuru-Guzik. Language models can learn complex molecular distributions. *Nature Communications*, 13(1):1–10, 2022.
- [11] Kohulan Rajan, Achim Zielesny, and Christoph Steinbeck. Decimer: towards deep learning for chemical image recognition. *Journal of Cheminformatics*, 12(1):1–9, 2020.
- [12] Kohulan Rajan, Achim Zielesny, and Christoph Steinbeck. Stout: Smiles to iupac names using neural machine translation. *Journal of Cheminformatics*, 13(1):1–14, 2021.

- [13] Nathan C Frey, Vijay Gadepally, and Bharath Ramsundar. Fastflows: Flow-based models for molecular graph generation. *arXiv preprint arXiv:2201.12419*, 2022.
- [14] Geemi P Wellawatte, Aditi Seshadri, and Andrew D White. Model agnostic generation of counterfactual explanations for molecules. *Chemical science*, 13(13):3697–3705, 2022.
- [15] Jan H Jensen. A graph-based genetic algorithm and generative model/monte carlo tree search for the exploration of chemical space. *Chemical science*, 10(12):3567–3572, 2019.
- [16] Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Junction tree variational autoencoder for molecular graph generation. In *International conference on machine learning*, pages 2323–2332. PMLR, 2018.
- [17] Yutong Xie, Chence Shi, Hao Zhou, Yuwei Yang, Weinan Zhang, Yong Yu, and Lei Li. Mars: Markov molecular sampling for multi-objective drug discovery. *arXiv preprint arXiv:2103.10432*, 2021.
- [18] Emmanuel Bengio, Moksh Jain, Maksym Korablyov, Doina Precup, and Yoshua Bengio. Flow network based generative models for non-iterative diverse candidate generation. *Advances in Neural Information Processing Systems*, 34:27381–27394, 2021.
- [19] Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Multi-objective molecule generation using interpretable substructures. In *International conference on machine learning*, pages 4849–4859. PMLR, 2020.
- [20] Soojung Yang, Doyeong Hwang, Seul Lee, Seongok Ryu, and Sung Ju Hwang. Hit and lead discovery with explorative rl and fragment-based molecule generation. *Advances in Neural Information Processing Systems*, 34:7924–7936, 2021.
- [21] Daniel Flam-Shepherd, Alexander Zhigalin, and Alán Aspuru-Guzik. Scalable fragment-based 3d molecular design with reinforcement learning. *arXiv preprint arXiv:2202.00658*, 2022.
- [22] Minghao Guo, Veronika Thost, Beichen Li, Payel Das, Jie Chen, and Wojciech Matusik. Data-efficient graph grammar learning for molecular generation. *arXiv preprint arXiv:2203.08031*, 2022.
- [23] Qi Liu, Miltiadis Allamanis, Marc Brockschmidt, and Alexander Gaunt. Constrained graph variational autoencoders for molecule design. *Advances in neural information processing systems*, 31, 2018.
- [24] Pavel Polishchuk. Crem: chemically reasonable mutations framework for structure generation. *Journal of Cheminformatics*, 12(1):1–18, 2020.
- [25] WJ Wiswesser. Simplified chemical coding for automatic sorting and printing machinery. Reading, Pa., Willson Products Inc, 1951.
- [26] William J Wiswesser. 107 years of line-formula notations (1861-1968). *Journal of Chemical Documentation*, 8(3):146–150, 1968.
- [27] John J Vollmer. Wiswesser line notation: an introduction. *Journal of Chemical Education*, 60(3):192, 1983.
- [28] H Winston Hayward. *A new sequential enumeration and line formula notation system for organic compounds*. Number 21. Office of Research and Development, Patent Office, 1961.
- [29] Herman Skolnik and Avis Clow. A notation system for indexing pesticides. *Journal of Chemical Documentation*, 4(4):221–227, 1964.
- [30] R Webster Homer, Jon Swanson, Robert J Jilek, Tad Hurst, and Robert D Clark. Sybyl line notation (sln): a single notation to represent chemical structures, queries, reactions, and virtual libraries. *Journal of chemical information and modeling*, 48(12):2294–2307, 2008.
- [31] Tianhong Zhang, Hongli Li, Hualin Xi, Robert V Stanton, and Sergio H Rotstein. Helm: a hierarchical notation language for complex biomolecule structure representation, 2012.

- [32] Diego Garay-Ruiz, Carles Bo, and Diego Garay Ruiz. Human-readable smiles: Translating cheminformatics to chemistry. 2021.
- [33] Xinhao Li and Denis Fourches. Smiles pair encoding: a data-driven substructure tokenization algorithm for deep learning. *Journal of Chemical Information and Modeling*, 61(4):1560–1569, 2021.
- [34] Products: Structure data downloads. URL <https://search.emolecules.com/info/products-data-downloads.html>.
- [35] Barbara Zdrzil and Rajarshi Guha. The rise and fall of a scaffold: a trend analysis of scaffolds in the medicinal chemistry literature. *Journal of Medicinal Chemistry*, 61(11):4688–4703, 2017.
- [36] Peter Ertl and Tim Schuhmann. A systematic cheminformatics analysis of functional groups occurring in natural products. *Journal of natural products*, 82(5):1258–1263, 2019.
- [37] Suliman Sharif, Ruibin Liu, Asuka A. Orr, Daniel Khavrutskii, Sunhwan Jo, Bettina Lier, Anastasia Croitoru, Chris Burke, Aziza Frank, Jacob Weiner, Nathaniel McClean, Aarion Romany, Mingtian Zhao, Takayuki Serizawa, Jared Deacon, Ian Jones, Shaoqi Zhan, Anmol Kumar, Mike Woster, Rebecca Pinette-Dorin, Elena Y. Chow, Sevien Schulhoff, and Alexander D. MacKerell Jr. Global-Chem: A Chemical Knowledge Graph of common small molecules and their IUPAC/SMILES/SMARTS for selection of compounds relevant to diverse chemical communities, 8 2022.
- [38] Jameed Hussain and Ceara Rea. Computationally efficient algorithm to identify matched molecular pairs (mmps) in large data sets. *Journal of chemical information and modeling*, 50(3):339–348, 2010.
- [39] Chunquan Sheng and Wannian Zhang. Fragment informatics and computational fragment-based drug design: an overview and update. *Medicinal Research Reviews*, 33(3):554–598, 2013.
- [40] Tairan Liu, Misagh Naderi, Chris Alvin, Supratik Mukhopadhyay, and Michal Brylinski. Break down in order to build up: decomposing small molecules for fragment-based drug design with e molfrac. *Journal of chemical information and modeling*, 57(4):627–631, 2017.
- [41] Peter S Kutchukian, Sung-Sau So, Christian Fischer, and Chris L Waller. Fragment library design: Using cheminformatics and expert chemists to fill gaps in existing fragment libraries. In *Fragment-Based Methods in Drug Discovery*, pages 43–53. Springer, 2015.
- [42] Simon Müller. Flexible heuristic algorithm for automatic molecule fragmentation: application to the unifac group contribution model. *Journal of cheminformatics*, 11(1):1–12, 2019.
- [43] Jörg Degen, Christof Wegscheid-Gerlach, Andrea Zaliani, and Matthias Rarey. On the art of compiling and using ‘drug-like’ chemical fragment spaces. *ChemMedChem: Chemistry Enabling Drug Discovery*, 3(10):1503–1507, 2008.
- [44] John J Irwin, Khanh G Tang, Jennifer Young, Chinzorig Dandarchuluun, Benjamin R Wong, Munkhzul Khurelbaatar, Yurii S Moroz, John Mayfield, and Roger A Sayle. Zinc20—a free ultralarge-scale chemical database for ligand discovery. *Journal of chemical information and modeling*, 60(12):6065–6073, 2020.
- [45] Peter Ertl and Ansgar Schuffenhauer. Estimation of synthetic accessibility score of drug-like molecules based on molecular complexity and fragment contributions. *Journal of cheminformatics*, 1(1):1–11, 2009.
- [46] G Richard Bickerton, Gaia V Paolini, Jérémy Besnard, Sorel Muresan, and Andrew L Hopkins. Quantifying the chemical beauty of drugs. *Nature chemistry*, 4(2):90–98, 2012.
- [47] Steven A Lopez, Benjamin Sanchez-Lengeling, Julio de Goes Soares, and Alán Aspuru-Guzik. Design principles and top non-fullerene acceptor candidates for organic photovoltaics. *Joule*, 1(4):857–870, 2017.

- [48] Daniil Polykovskiy, Alexander Zhebrak, Benjamin Sanchez-Lengeling, Sergey Golovanov, Oktai Tatanov, Stanislav Belyaev, Rauf Kurbanov, Aleksey Artamonov, Vladimir Aladinskiy, Mark Veselov, Artur Kadurin, Simon Johansson, Hongming Chen, Sergey Nikolenko, Alan Aspuru-Guzik, and Alex Zhavoronkov. Molecular Sets (MOSES): A Benchmarking Platform for Molecular Generation Models. *Frontiers in Pharmacology*, 2020.
- [49] Mark A Murcko and Guy W Bemis. The properties of known drugs. 1. *Molecular Frameworks. J. Med. Chem.*, 39:2887–2893, 1996.
- [50] Kristina Preuer, Philipp Renz, Thomas Unterthiner, Sepp Hochreiter, and Günter Klambauer. Fréchet ChemNet distance: A metric for generative models for molecules in drug discovery. 58 (9):1736–1741. ISSN 1549-9596. doi: 10.1021/acs.jcim.8b00234. URL <https://doi.org/10.1021/acs.jcim.8b00234>. Publisher: American Chemical Society.
- [51] Mustafa Hajj, Kyle Istvan, and Ghada Zamzmi. Cell complex neural networks. *arXiv preprint arXiv:2010.00743*, 2020.