

Learning from flowsheets

A generative transformer model for autocompletion of flowsheets

Vogel, Gabriel; Schulze Balhorn, Lukas; Schweidtmann, Artur M.

DOI

[10.1016/j.compchemeng.2023.108162](https://doi.org/10.1016/j.compchemeng.2023.108162)

Publication date

2023

Document Version

Final published version

Published in

Computers and Chemical Engineering

Citation (APA)

Vogel, G., Schulze Balhorn, L., & Schweidtmann, A. M. (2023). Learning from flowsheets: A generative transformer model for autocompletion of flowsheets. *Computers and Chemical Engineering*, 171, Article 108162. <https://doi.org/10.1016/j.compchemeng.2023.108162>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.



Learning from flowsheets: A generative transformer model for autocompletion of flowsheets

Gabriel Vogel, Lukas Schulze Balhorn, Artur M. Schweidtmann*

Delft University of Technology, Department of Chemical Engineering, Process Intelligence Research Group, Van der Maasweg 9, Delft, 2629 HZ, Netherlands

ARTICLE INFO

Keywords:

Flowsheet synthesis
Flowsheet completion
SFILES 2.0
Natural language processing
Generative transformer model

ABSTRACT

We propose a novel method enabling autocompletion of chemical flowsheets. This idea is inspired by the autocompletion of text. We represent flowsheets as strings using the text-based SFILES 2.0 notation and learn the grammatical structure of the SFILES 2.0 language and common patterns in flowsheets using a transformer-based language model. We pre-train our model on synthetically generated flowsheet topologies to learn the flowsheet language grammar. Then, we fine-tune our model in a transfer learning step on real flowsheet topologies. Finally, we use the trained model for causal language modeling to autocomplete flowsheets. Eventually, the proposed method can provide chemical engineers with recommendations during interactive flowsheet synthesis. The results demonstrate a high potential of this approach for future AI-assisted process synthesis but also reveal the limitations at the present state and the next steps that need to be taken to deploy this technique in realistic flowsheet synthesis scenarios.

1. Introduction

One important step in early process synthesis is the design of the flowsheet topology consisting of the selection and arrangement of unit operations and their connections. In an iterative procedure, engineers mostly perform this step manually based on experience, common design heuristics, and optimization-based methods (Chen and Grossmann, 2017; Biegler et al., 1997).

Recent years have revealed a very high potential for artificial intelligence (AI) in various disciplines of natural science and engineering (e.g., Jumper et al., 2021; Silver et al., 2016; Secinaro et al., 2021). Advances in natural language processing (NLP), a sub-field of AI, support several applications in daily life, such as intelligent text completion in search engines or email programs. Recently, NLP methods have also been successfully applied to other domains. Github copilot uses text generation models to autocomplete computer code supporting software developers (Chen et al., 2021). Another example are NLP applications in the molecular world using the text-based Simplified Molecule Input-Line Entry-System (SMILES) (Weininger et al., 1989) notation which is a common representation format for molecules and chemical reactions. Recurrent neural networks (RNNs) trained on SMILES strings have been used as generative models in de novo drug design (Segler et al., 2017). Recently, transformer-based language models (Vaswani et al., 2017) have been applied to several tasks, such as the prediction of molecular properties (Chithrananda et al., 2020), chemical reaction

prediction with the Molecular Transformer (Schwaller et al., 2019), and the prediction of retrosynthetic pathways (Schwaller et al., 2020).

To this date, there have been only few previous works that use AI for process synthesis (Schweidtmann et al., 2021; Lee et al., 2018; Venkatasubramanian, 2018). Notably, there have been many previous works on surrogate modeling and (superstructure) optimization (Henao and Maravelias, 2011; Kim and Boukouvala, 2020; Schweidtmann and Mitsos, 2019; Huster et al., 2020) while only a few works actually create process topologies. Recently, there have been advances toward using Reinforcement Learning (RL) approaches for process synthesis (Khan and Lapkin, 2020; Göttl et al., 2021b,a; Midgley, 2020; Stops et al., 2022). RL does not utilize data from existing flowsheets but strives to train an agent to learn the task of process synthesis. Contrary, the variety and large amount of historical chemical processes and their corresponding flowsheets pose many opportunities for data-driven models in the field of process synthesis (Zhang et al., 2018; Zheng et al., 2022). However, unleashing the potential of data-driven AI in the chemical engineering domain also comes with several difficulties (Schweidtmann et al., 2021; Weber et al., 2021), such as the design of meaningful information representations allowing advanced AI methods to understand complex data. In parallel to the developments of our proposed work, Oeing et al. (2022) very recently published a data-driven approach for AI-assisted development of piping and instrumentation diagrams (P&IDs). The authors collected 35 P&IDs in

* Corresponding author.

E-mail address: a.schweidtmann@tudelft.nl (A.M. Schweidtmann).

DEXPI format and converted them to graphs. Then, they generated a training data set of linear sequences of six sequential unit operations by sampling random walks from the graphs. Oeing et al. (2022) trained an RNN on the sequential data to predict subsequent equipment. Furthermore, they used the graph data set to train a graph neural network (GNN) to perform consistency checks during the drawing of P&IDs, by comparing newly drawn equipment with patterns in the data set.

We propose a novel methodology enabling autocompletion of chemical process flowsheets by learning the structure and patterns from a dataset of process flow diagrams (PFDs) using a transformer-based NLP model. The underlying idea of our approach is to apply causal language modeling to autocomplete flowsheets comparable to sentence completion or, in general, text generation of human language. We mine and digitize PFDs from literature and also generate an additional synthetic PFD dataset. In order to train and deploy the transformer model on complex flowsheet topologies, we represent PFDs using our recent extension of the text-based Simplified Flowsheet Input-Line Entry-System (SFILES) 2.0 (d'Anterrosches, 2006; Vogel et al., 2022b). Eventually, the proposed flowsheet-completion approach aims to provide engineers with recommendations for the flowsheet structure during AI-assisted iterative process synthesis. In comparison to the work by Oeing et al. (2022), we train a transformer model that learns from whole flowsheets rather than training an RNN on linear sequences. In particular, the SFILES 2.0 format allows us to represent and learn from complex flowsheet topologies including recycles and branchings (e.g., in distillation columns) while the previous work is limited to linear sequences. In addition, our model has a maximum input length of 512 unit operations per flowsheet compared to 6 units in the previous work. Likewise, our approach also autocompletes complete flowsheets (incl. complex recycles and branching) compared to the prediction of the next unit operation. Finally, we train our model on a comprehensive dataset of over 8000 PFDs compared to 35 P&IDs.

The remainder of this paper is structured as follows: In Section 2, we briefly explain the transformer architecture and the used flowsheet representations. Section 3 explains all components of the flowsheet completion model, called Generative Flowsheet Transformer. Afterward, the used data for pre-training and fine-tuning of the model as well as the training procedure and results are described in Sections 4 and 5. In Section 6, we demonstrate the autocompletion technique using illustrative examples. In Section 6.3 we also point out the current limitations of the proposed model regarding applicability in a context of realistic design problems and discuss the next essential steps that need to be taken to develop a more intelligent model.

2. Background

This section briefly explains the concepts of the original transformer model architecture (Vaswani et al., 2017) and further the typical modified model architecture for text generation applications. Afterward, it recaps the used flowsheet representations, namely flowsheet graphs and the SFILES 2.0 notation. The latter is used to represent the flowsheet data in a text-based manner and to use NLP models on the flowsheet data.

2.1. Original transformer architecture

For many years, the state-of-the-art in the field of NLP had been recurrent architectures using LSTMs (Hochreiter and Schmidhuber, 1997) or GRUs (Cho et al., 2014). In the recent past, Vaswani et al. (2017) published the novel transformer model architecture using self-attention which is currently facing a hype in the AI community. It does not use recurrence but processes the entire sequence at once and therefore allows for significantly more parallelization in comparison to RNNs. Furthermore, self-attention layers are faster than recurrent layers in terms of computational complexity (Vaswani et al., 2017). The reduced

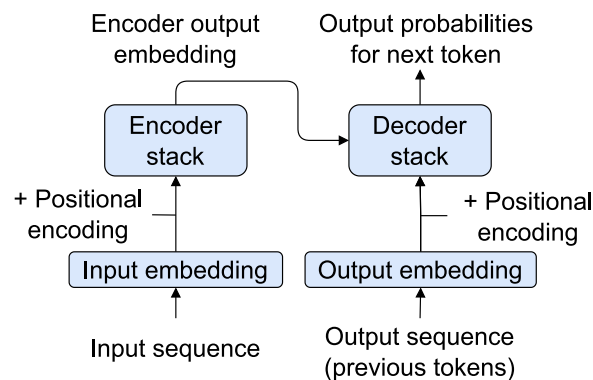


Fig. 1. Simplified illustration of transformer architecture derived from Vaswani et al. (2017) consisting of encoder and decoder stack.

computational cost is one major reason for the growing popularity not only in the field of NLP but also in the whole AI community. Transformer models have demonstrated breakthrough performances in a variety of applications (Vaswani et al., 2017; Schwaller et al., 2019; Lakew et al., 2018) and have, in many cases, the potential to replace RNNs and convolutional neural networks.

The original transformer architecture (Vaswani et al., 2017) is a neural sequence translation model which is usually used to translate an input sequence to an output sequence, e.g. a German to an English sentence. It consists of an encoder stack and a decoder stack, as shown in the simplified illustration in Fig. 1. The encoder converts a sequence of words or characters (input tokens) into a numerical representation that can be further processed in the model. It captures each token and its meaning in the sequence. This numerical representation is then fed to the decoder stack, which converts it back to a sequence of words or characters (generated output tokens) in the target language. The generation of the translated sequence is a token by token generation. This means, that the decoder repeatedly uses the encoder's numerical output and previously generated output tokens to compute the probabilities for the next output token.

In the original architecture, the encoder stack and decoder stack each comprise $N = 6$ identical layers. Multiple layers are used to enhance the model's ability to capture the whole complexity of a language. Each encoder layer contains two sub-layers with subsequent layer normalization. Each decoder layer contains three sub-layers with subsequent layer normalization. There are two general types of sub-layers used. First, the attention sub-layers comprise mathematical operations that help understand the language and the interdependence of tokens in a language. Second, the position-wise feed-forward sub-layers further process the attention output before feeding it to the next encoder or decoder layer.

Since recurrent components are completely removed in the transformer architecture, before input and output embeddings are passed to the encoder and decoder, respectively, positional encoding is applied (see Fig. 1). Positional encoding ensures that the information of the order of tokens in the sequence is taken into account. In the original implementation (Vaswani et al., 2017), sine and cosine functions of different frequencies depending on the token position are added to the embeddings allowing the model to know each token position in the sequence.

One important core component of the transformer architecture are the attention sub-layers. The calculation of attention takes a query vector q , key vector k , and value vector v for each input token and compares all queries against all keys resulting in scores for query-key compatibility. The compatibility scores are then used as weights to calculate the attention output as a weighted sum of the values. In practice, the attention is computed for all inputs of an input sequence in parallel, putting together all query, key, and value vectors in the




Fig. 2. Simplified illustration of a decoder-only architecture for auto-regressive text-generation.

query matrix Q , key matrix K , and value matrix V . This finally yields a matrix as attention output. In the transformer (Vaswani et al., 2017), the scaled dot-product attention is implemented, which is calculated from the queries Q , keys K , and values V as

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V, \quad (1)$$

where d_k is the dimension of the key vectors. The transformer in specific uses multi-head attention, consisting of several scaled dot-product attention layers in parallel. For the multi-head attention, the keys, values, and queries are linearly projected into h learned projections of the dimension n_{embd}/h , enabling multiple heads to learn from different representation subspaces in parallel (Vaswani et al., 2017).

In the original architecture, multi-head attention is used as self-attention layers in the encoder, masked self-attention in the decoder, and encoder–decoder attention to combine the vector embedding of the encoder with the previous decoder outputs. Hereby, self-attention means that query, key, and value matrices are calculated from the same input sequence. Therefore, the computed self-attention represents each token and its meaning in the sequence. Self-attention in the encoder considers both the left and right context of each token (bidirectional). Contrary, in the case of masked self-attention in the decoder, only the left context is used, meaning that subsequent positions of each token are masked out (unidirectional). In practice, a mask is added to the matrix multiplication product of Q and K in Eq. (1), i.e., by setting the softmax input of the respective positions to $-\infty$ and thus the respective outputs scores of the softmax to zero. This way the auto-regressive property of the decoder is preserved (Vaswani et al., 2017). The encoder–decoder attention uses the encoder’s output as keys and values and the decoder’s masked self-attention output as queries.

2.2. Auto-regressive transformer for text generation

Besides the original architecture, depending on the application, encoder and decoder stacks can be modified or left out entirely. Auto-regressive models for text generation, also called causal language modeling, typically only use the decoder part of the original transformer, as shown in Fig. 2. Auto-regressive means that previously generated tokens are added to the input sequence for the next token generation. More specifically, at each time step, the decoder-only transformer model outputs probabilities for different tokens suitable as the next token in the sequence. The selected token is then added to the previous input sequence (dashed line in Fig. 2) before the decoder computes the following outputs.

In recent years, many model architectures for text generation and pre-trained models were published, such as GPT-2 (Radford et al., 2018), Transformer-XL (Dai et al., 2019), and XLNet (Yang et al.).

2.3. Graph- and text-based representation of flowsheets

This section briefly summarizes the flowsheet representations that we used in the data processing steps. The data sets, described in Section 4, are first created in graph format and later converted to the text-based SFILES 2.0 strings (d’Anterrosches, 2006; Vogel et al., 2022b).

Flowsheets can be represented as directed heterogeneous graphs (Vogel et al., 2022b; Zhang et al., 2018) with unit operations as nodes and stream connections between the unit operations as directed edges. Fig. 3 shows an example flowsheet with a two-stream heat exchanger, reactor, distillation system, and a recycle loop. This flowsheet can be represented as the flowsheet graph shown in Fig. 4, whereby the heat exchanger unit operation is divided into two nodes, one for each stream compartment.

Using the SFILES 2.0 notation (Vogel et al., 2022b) and our open-source conversion implementation, provided in our [GitHub repository](#) (Vogel et al., 2022a) yields the corresponding SFILES 2.0 string

```
(raw) (hex) {1} (r) <&| (raw) (pp) &| (mix) <1 (v) (dist) [{
  tout} (prod)] {bout} (spl) 1 (prod) n | (raw) (hex)
{1} (prod)
```

for the flowsheet graph in Fig. 4. In an SFILES 2.0 string, unit operations are represented in parentheses. Two consecutive unit operations in parentheses imply a directed stream connection from the left to the right unit operation. In the case of branching in the process such as after a distillation column (dist) (one input and two output streams), the branches are noted in brackets. The brackets are omitted for the branch that is noted last in the SFILES 2.0 string (in the given example the branch continuing with (spl)). Recycles are noted using numbers # to reference the recycle start node (here (spl)) and <# to reference the recycle end node (here (mix)). Furthermore, tags in braces are used to indicate whether the branch is a top or bottom product. In the case of converging branches, the second branch is inserted in the string, surrounded by <&| and &| (here <&| (raw) (pp) &|). Multi-stream heat exchangers are separated in one node per stream compartment and marked with a number in braces, capturing which streams are heat-integrated.

3. Generative flowsheet transformer

This section first provides an overview of the general procedure from an incomplete to a completed flowsheet in Section 3.1. Then, we specify details of the single components of the Generative Flowsheet Transformer, including the SFILES 2.0 tokenization (Section 3.2), the decoder stack architecture (Section 3.3), and the decoding strategies for flowsheet completion (Section 3.4).

3.1. Overview

The proposed flowsheet completion methodology is illustrated in Fig. 5(a) and described in the following. In step 1, the incomplete flowsheet graph is converted to the corresponding SFILES 2.0 string, as described in Section 2.3. Afterward, in step 2, the string is tokenized using the SFILES 2.0 tokenizer (Section 3.2). The resulting input embedding is passed to the decoder stack (Section 3.3), which computes the output probabilities for the next token prediction. Step 4 comprises the selection of the next token from the decoder output, which is determined by the decoding strategy, as described in Section 3.4. After the SFILES 2.0 string completion is finished by reaching a defined end token or by controlling the number of tokens to be generated (step 5), the string is converted back to the autocompleted flowsheet graph (step 6). This procedure could be embedded in process simulation software for the interactive autocompletion of flowsheets.

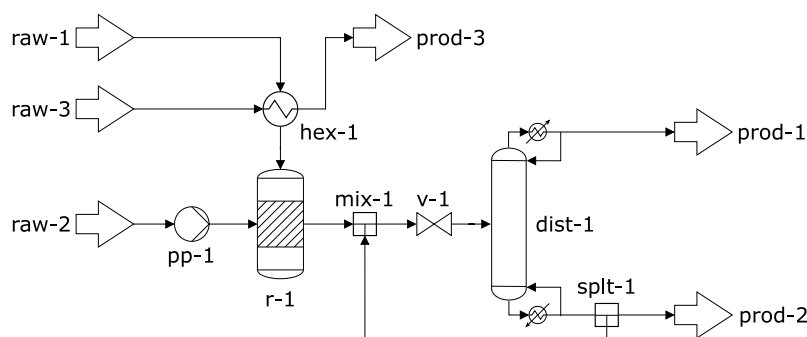


Fig. 3. Simple chemical process flowsheet with branchings, recycle stream, and different mass trains.




Fig. 4. Graph representation of flowsheet in Fig. 3.

3.2. Tokenizer

Tokenization is the first text processing step in NLP. Therein, text input is usually processed as a sequence of tokens, whereby the tokens are either words or other chunks of the input sequence. Common strategies to identify tokens in natural language are word-level, WordPiece (Wu et al., 2016), and most recently, Byte Pair Encoding (BPE) tokenization (Sennrich et al., 2015). After tokenizing the input sequence, each token is converted to a vector of size n_{embd} , called numerical embedding. This vector is calculated by multiplying a one-hot-encoded vector of the token with a weight matrix, learned by the language model during the training procedure. Combining the input vectors of a sequence of tokens yields a matrix that is called the input embedding.

We trained our language model on the SFILES 2.0 language, which significantly differs from human language. Therefore, common tokenization strategies such as word-level, WordPiece, and Byte-pair encoding are not suitable for our task. Thus, we developed a tailored tokenizer for SFILES 2.0 which identifies, for instance, a unit operation (e.g. (hex)), a branching, stream tags (e.g. {tout}), a recycle connection number, etc., using a regular expression. This approach is inspired by the tokenizer for the molecular representation using SMILES (Schwaller et al., 2018). Our proposed regular expression is

```
r"(\(. *?\)|\{. *?\}|\%([0-9]{3})\|
  |\%[0-9]{2}\| |\[[\<.?[0-9]\<&\|\| (?<!\<
  \&\| \|n\| \| (?<!\&) (?<!\n)\| \| \& (?!\|)
  |\[/[0-9] | [0-9])".
```

The following example shows an SFILES 2.0 string before and after the tokenization.

Before: (raw)(hex)(r)(mix)<1(v)(dist)[{tout}(prod)]{bout}(spl-1)(prod)

After: (raw)(hex)(r)(mix)<1(v)(dist)[{tout}(prod)]{bout}(spl-1)(prod)

3.3. Decoder-only architecture for causal language modeling

Since we aim to complete SFILES 2.0 strings using causal language modeling, our proposed architecture is based on an auto-regressive decoder-only transformer model as described in Section 2.2. The SFILES 2.0 language with a small vocabulary size of less than 100 is rather simple compared to human language with a vocabulary size of approximately 500k words in Finnish or approximately 90k words in English (Birch et al., 2008). Thus, we use the architecture of the small version of the GPT-2 transformer model (Radford et al., 2018), which contains a decoder stack consisting of $N = 12$ decoder layers. Each decoder layer comprises a masked multi-head self-attention sub-layer with $h = 12$ attention heads and a feed-forward sub-layer. With an embedding dimension of $n_{embd} = 768$ the number of parameters in this architecture adds up to 85.9 million. The high number of parameters increases the capacity of the model to capture the entire complexity of a language. On the other hand, this also increases the risk of overfitting the model on our training data. The GPT-2 configuration we used, by default implements dropout during the training procedure in the embedding, attention, and fully-connected layers to counteract overfitting. Furthermore, we make use of a technique called early-stopping during the training procedure, which monitors the error on a validation set and stops the training once this error does not decrease anymore. In future works, a more rigorous investigation of the optimal model size and hyperparameters could be performed using Bayesian optimization or information criteria.

3.4. Decoding strategies

When completing a flowsheet, as shown in Fig. 5, the input is either an empty flowsheet represented as an empty input sequence or an incomplete flowsheet represented as an incomplete SFILES 2.0 string. In this work, we use open-ended text generation to complete SFILES 2.0 strings. Starting from an input sequence of m tokens t_1, \dots, t_m , considered the context for the language generation, the trained model generates n tokens in an auto-regressive manner until the sequence

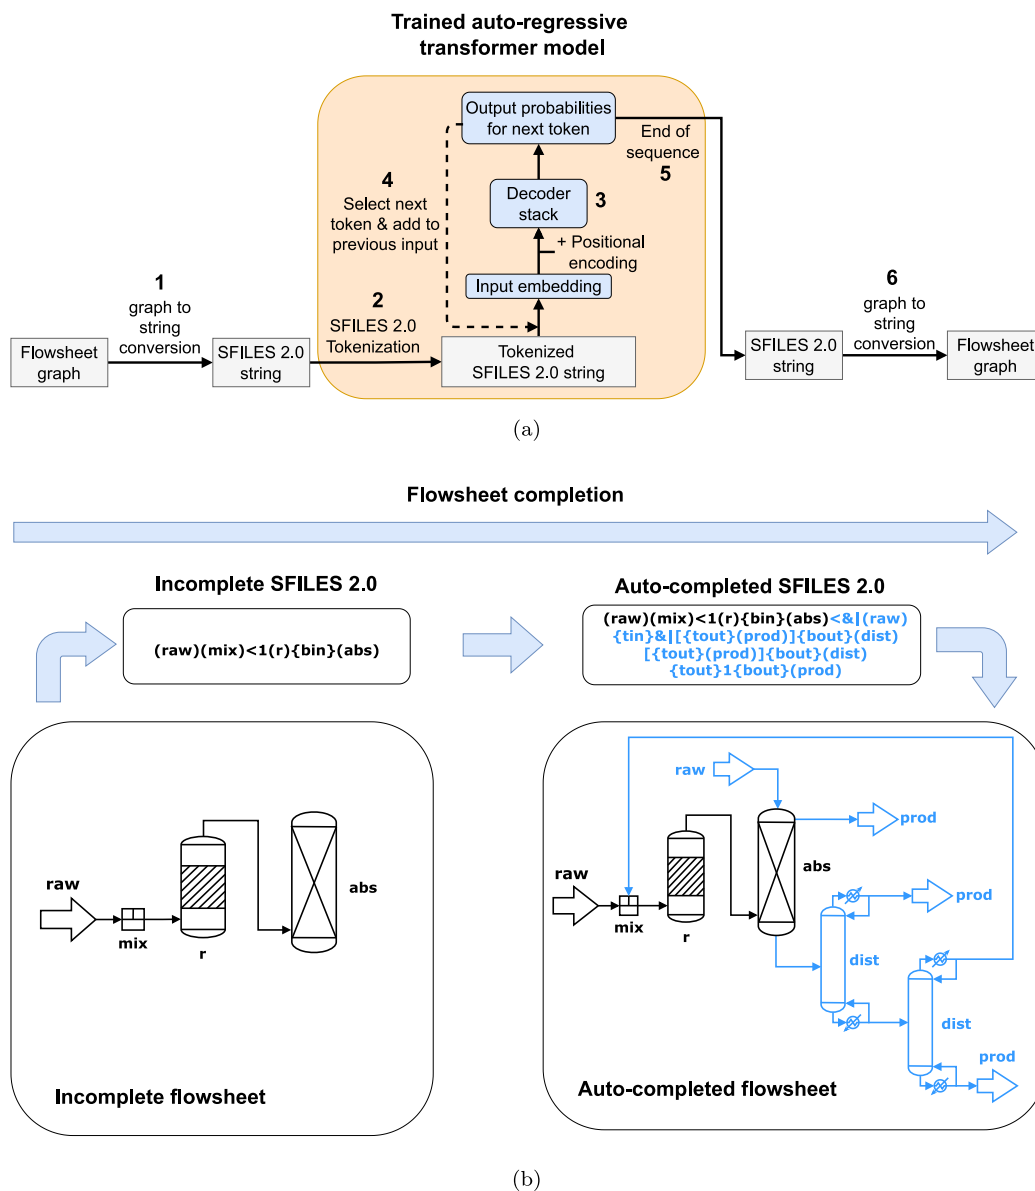


Fig. 5. Overview of flowsheet completion with the Generative Flowsheet Transformer (a) Incomplete flowsheet graph is converted to incomplete SFILES 2.0 string (1). Auto-regressive transformer model completes string (2,3,4,5). autocompleted SFILES 2.0 string is converted to completed flowsheet graph (6). (b) Example auto-completion of a flowsheet.

is completed (Holtzman et al., 2020) (see Section 3.3). With the assumption that the probability distribution is composed of the product of conditional next token distributions (Holtzman et al., 2020), the final probability of the sequence is given as

$$P(t_{1:m+n}) = \prod_{i=1}^{m+n} P(t_i | t_1, \dots, t_{i-1}). \quad (2)$$

Different decoding (or text generation) strategies define the selection of the next token at each step. Common strategies are greedy search, beam search, top- k sampling, and top- p sampling.

Greedy search simply selects the token with the highest probability at each time step. This increases the chance of missing high-probability tokens hidden after low-probability tokens. Therefore, greedy search often does not lead to sequences with a maximized final probability.

Beam search aims to overcome the weakness of greedy search, i.e. by finding the optimal sequence by maximizing the final probability $P(t_{1:m+n})$ according to Eq. (2). Constructing a tree of the possible tokens at each time step, the beam search is similar to a breadth-first search through the tree to find the path (sequence) with the highest likelihood.

The only difference is that beam search takes a parameter, beam-width w_{beam} , which determines the number of considered tokens at each step. However, several studies found that beam search often yields repetitive text (Shao et al., 2017; Vijayakumar et al., 2016) and lower-quality text compared to other decoding strategies that include sampling (Fan et al., 2018; Holtzman et al., 2020).

Top- k and top- p sampling increase the chance of also selecting lower-probability tokens which often yields more diversity of the generated sequences. The parameters k and p truncate the distribution of token probabilities and prevent the model from randomly sampling from the whole probability distribution but rather from a limited selection of more-likely tokens.

The top- p sampling decoding strategy addresses the problem of fixed k values in top- k sampling. Its key idea is to use the shape of the probability distribution to define a confidence region at each time step (Holtzman et al., 2020). The number of candidates in this confidence region dynamically changes and thus eliminates the problem of selecting a fixed number k of candidates, as in top- k sampling. With the definition of a threshold p , the smallest set of top- p candidates

$C^{(p)}$ is selected, such that the cumulative probability mass exceeds the threshold. Eq. (3) shows this relation.

$$\sum_{t \in C^{(p)}} P(t|t_{1:i-1}) \geq p \quad (3)$$

Depending on the shape of the probability distribution, the number of candidates within that set changes. Usually, high values of p (between 0.9 and 1) are selected such that the set $C^{(p)}$ will take up most of the probability mass representing the nucleus (Holtzman et al., 2020). The probability distribution is then re-scaled only using the set $C^{(p)}$ and the next token is sampled from the new distribution (Holtzman et al., 2020).

We use beam search with a fixed beam-width of $w_{beam} = 5$. The selection of w_{beam} is a compromise of getting better results for long generated sequences while keeping the computational cost low. To introduce more diversity in the generated sequences, we choose top- p sampling as our second decoding strategy with a cumulative probability threshold of $p = 0.9$.

4. Data

This section explains the data sets that we use for the training of the Generative Flowsheet Transformer. The pre-training data set comprises synthetically generated flowsheets, while the fine-tuning data set consists of real flowsheets extracted from chemical process simulation files.

4.1. Generated data for pre-training

Language models are complex architectures with a high number of trainable weights. When dealing with human language, training these weights usually requires several Gigabytes to Terrabytes of text (Radford et al., 2018; Brown et al., 2020). The SFILES 2.0 language is less complex than human language, consisting of a small vocabulary. Still, for the model to learn the SFILES 2.0 grammar of flowsheet topologies, we need a reasonable amount of data. There is currently no (public) database or data set of flowsheets (Schweidtmann, 2022). Hence, we construct a synthetic data set of realistic flowsheet graphs for pre-training of our model.

The synthetic flowsheet generation builds up flowsheet graphs by a drawing random samples from a Markov chain-like process based on known flowsheet design heuristics. We subdivide flowsheets into the following sub-process categories: Initialization with feed(s), reaction, thermal separation (distillation, rectification), countercurrent separation (absorption, extraction), filtration (gas, liquid), centrifugation, and purification as the last subprocess. As illustrated in Fig. 6, after initializing the flowsheet graph with raw materials, including feed pre-processing, the selection of the first sub-process, excluding purification, is a Markov transition with fixed probabilities where transition probabilities do not depend on previous unit operations. Within each sub-process, we further sample from a set of patterns specifying how the inlet and outlet streams are processed, e.g., with additional temperature or pressure change unit operations. Also, we include design heuristics such as adding recycles, performing heat integration in the reaction sub-process, or adding reactants (Zhang et al., 2018). In general, the sub-processes lead to several outlet streams, in the following referred to as branches. For each branch, we transition to the “Next sub-process” state, followed by a Markov transition to the next sub-process. This selection differs from the first sub-process selection by the additional purification sub-process. Note that once a branch reaches the purification step, it is determined to end as a product. After each branch ended in the purification step, the flowsheet graph generation is complete. Table 1 summarizes the properties of the generated data set after removing duplicates and flowsheet graphs with more than 50 nodes. The SFILES 2.0-based data set for pre-training is obtained by automatic conversion of the generated graphs to SFILES 2.0 strings (Vogel et al., 2022b).

Table 1

Properties of the used data sets and the number of samples in training (tr), validation (val), and test (te) data set.

	Generated data set	Real data set
samples _{tr}	6362	178
samples _{val}	796	23
samples _{te}	795	22
\bar{n}_{nodes}	19	25
$\sigma(n_{nodes})$	10	41
Vocabulary size	53	89

4.2. Real flowsheet data for fine-tuning

We collected 223 Aspen and DWSIM chemical process simulation files from the public domain. We automatically extract information from the simulation files and convert them to flowsheet graphs using our in-house software. Then, the process graphs are converted to SFILES 2.0 strings using our open-source code (Vogel et al., 2022b). Finally, we use the extracted flowsheet data for fine-tuning the pre-trained model.

Table 1 summarizes some key statistics of the synthetic and real flowsheet data set. The real data set shows a higher average node number, standard deviation, and vocabulary size compared to the generated data.

5. Training results and discussion

The following section first describes and discusses the results of the transformer model training, including pre-training and fine-tuning. Furthermore, the model is evaluated based on perplexity, a metric commonly used in NLP.

5.1. Model training

For pre-training and fine-tuning, we divide the dataset into a training (80%), validation (10%), and test set (10%). Table 1 shows the training, validation, and test set sizes. During pre-training, we used a batch size of 8, while during fine-tuning on the small data set, we used a batch size of 2 to update weights more frequently. The validation set was evaluated every 200 steps during pre-training and every 20 steps during fine-tuning. In both training procedures, we use early stopping with to prevent overfitting of the model to the training set, with the early stopping patience set to 3 during pre-training and to 20 during fine-tuning.

Fig. 7(a) shows the loss curves of the pre-training. The pre-training results show only small fluctuations in the loss curves and the gap between train and validation loss is small. This demonstrates a small generalization error on the synthetic dataset. This is likely due to the limited variability in the pre-training data, which is generated synthetically. Essentially, the data sets are drawn from the same probability distribution. Thus, the learned flowsheet patterns from the training set are representative for those in the synthetic validation and test set.

Fig. 7(b) shows the loss curves of the fine-tuning. In contrast to the pre-training, the fine-tuning shows considerable fluctuations of the loss curves and a larger gap between train and validation loss. The reason for this learning behavior is likely caused by the real flowsheet data. The real data consist of fewer flowsheets with higher variations in the number of nodes per flowsheet and an extended vocabulary size, i.e., more unit operation categories, as shown in Table 1. This leads to heavier fluctuations and a larger generalization error. In the future, we envision to create a public knowledge graph for flowsheets that will mitigate this issue (Schweidtmann, 2022).




Fig. 6. Random flowsheet graph generation scheme.

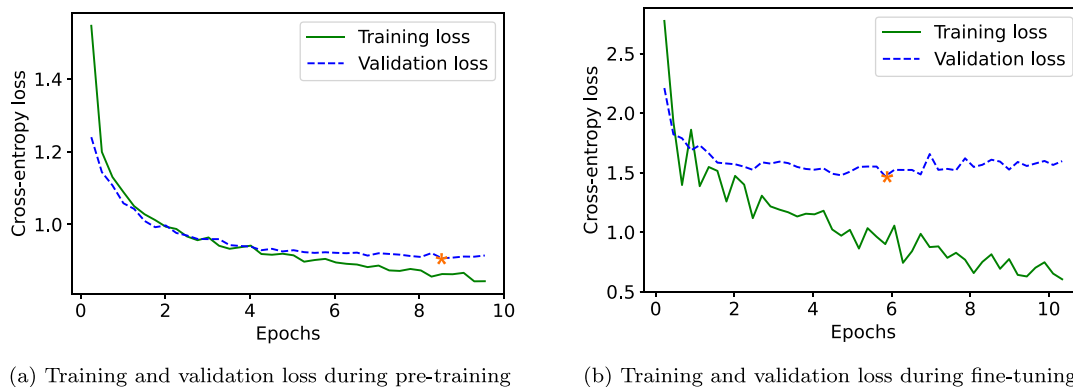


Fig. 7. Training and validation cross-entropy loss over number of training epochs. The asterisk marks the lowest validation loss during training which is used in the early stopping strategy to load the model parameters.

5.2. Perplexities

In the following, we evaluate the model results using the perplexity PP , which is the most common metric for the evaluation of text generation models. It measures the ability of a probability model to predict a sequence of unseen test data and is defined as the exponential of the negative average log-likelihood of a sequence, which is equivalent to the exponential of the cross-entropy loss obtained during the model training. For the sequence $T = (t_1, \dots, t_n)$ of n tokens, we calculate the perplexity as

$$PP(T) = \exp\left(-\frac{1}{n} \sum_i^n \log P(t_i | t_{1:i-1})\right).$$

In this equation, $P(t_i | t_{1:i-1})$ is the predicted conditional probability for token t_i given the tokens t_1, \dots, t_{i-1} . Low perplexities indicate that

Table 2

Perplexities of pre-trained and fine-tuned model for all used data sets.

Model	Data set	PP_{tr}	PP_{val}	PP_{te}
Pre-trained	Generated data	2.28	2.48	2.47
	Real data	41.61	64.16	25.91
Fine-tuned	Generated data	5.04	5.27	5.33
	Real data	2.17	3.98	4.75

the model does well in predicting the next tokens in the test data, while high perplexities imply that the learned probability distribution is different from the probability distribution of the test data.

Table 2 shows the perplexities of the pre-trained and fine-tuned models using both data sets. The pre-trained model shows similar perplexities on the generated train-, validation-, and test-set, because

the data is all drawn from the same probability distributions of unit operations and patterns. This is consistent with the small generalization error, observed in Section 5.1.

Pre-training with synthetic data is necessary but not sufficient for predicting real flowsheet data. Table 2 shows that the pre-trained model performs poorly on the data set of real flowsheets. Thus, we conclude that the generated data does not represent real flowsheet topologies sufficiently and fine-tuning on real data is necessary. Note that we also attempted to train a model from scratch only using the real data. Nevertheless, it resulted in worse training, validation, and test loss than our proposed approach. This is expected as transformer models are known to be very data hungry. We conclude that the pre-training on a larger corpus of SFILES 2.0 enables the model to learn the basic grammatical structure of the SFILES 2.0 notation. The learned knowledge enhances the model's performance when fine-tuning it on the real data.

The fine-tuning significantly improves the model performance on real data. After fine-tuning, the model performance on unseen data increases from 25.91 to 4.75. This is remarkable as the final performance on the real flowsheet data is in the same order of magnitude as the pre-trained model on the generated dataset (2.47). The results show that the model has the potential to learn flowsheet patterns from real flowsheets. Notably, the fine-tuned model performs slightly worse on the generated data (5.33) than the pre-trained model (2.47), which is expected as the model replaced some of the learned patterns of the generated data set with the patterns contained in real flowsheets.

6. Illustrative examples and discussion

In this section, we illustrate and compare the proposed autocompletion techniques on a few examples. We further discuss the applicability of the results and point out the limitations that our model faces at the present state.

6.1. Example 1

In the first, rather simple example, we start with one raw material and a heat exchanger as the input sequence. The start sequence represented as an SFILES 2.0 string is (raw)(hex). First, we proceed with beam search as the decoding strategy to complete the SFILES 2.0 string (blue is the generated text). We take the three sequences with the highest sequence probabilities according to Eq. (2). Second, we generate three examples, using top- p sampling to complete the same input sequence.

Start: (raw)(hex)

Beam search (descending sequence probability):

1. (raw)(hex)(r)[(prod)](hex)(flash)[(tout)(prod)](bout)(prod)
2. (raw)(hex)(r)(hex)(dist)[(tout)(hex)(prod)](bout)(hex)(prod)
3. (raw)(hex)(r)(hex)(flash)[(tout)(hex)(prod)](bout)(hex)(prod)

Top- p -sampling:

1. (raw)(hex)(mix)<1<&(r)&(dist)[(tout)(prod)](bout)(dist)(bout)1(tout)(prod)
2. (raw)(hex)(r)(hex)(dist)[(tout)(hex)(prod)](tout)(hex)(prod)
3. (raw)(hex)(dist)[(tout)(prod)](bout)(hex)(prod)

The sequences generated with beam search are grammatically correct SFILES 2.0 strings that can be converted to the flowsheet graphs shown in Figs. 8(a), 8(b), and 8(c). The three generated sequences with top- p sampling are also grammatically correct and can be converted to flowsheet graphs, shown in Figs. 9(a), 9(b), and 9(c). The order of

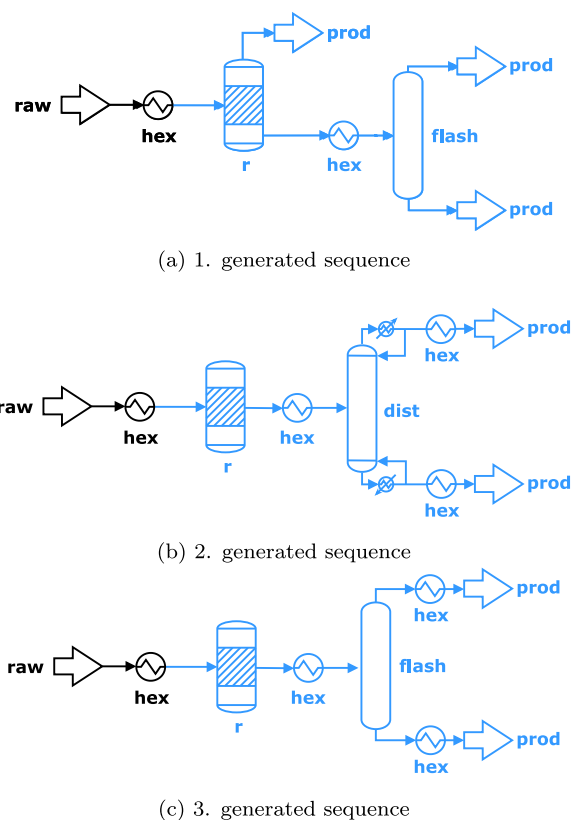


Fig. 8. Example 1: Completed flowsheets using beam search.

the generated sequences with top- p sampling is random and does not indicate an order of the sequence probabilities.

The example aims to show that the Generative Flowsheet Transformer is able to build a realistic flowsheet starting from an almost empty flowsheet graph. Beam search maximizes the final probability of the sequence, so in this case the generated flowsheets in Fig. 8 result in very short flowsheets with unit operation sequences (reactor, heat exchanger and flash/distillation system) commonly found in many chemical processes. Using top- p sampling can lead to a lower final probability of the generated sequences, thus, less common topologies such as in Fig. 9(a). Nevertheless, both selected decoding strategies lead to correct but relatively simple final flowsheets, engineers can quickly design in practice. A more realistic scenario would be the autocompletion of more complex flowsheets (e.g. in Section 6.1) or an interactive approach between the Generative Flowsheet Transformer and a flowsheet designer.

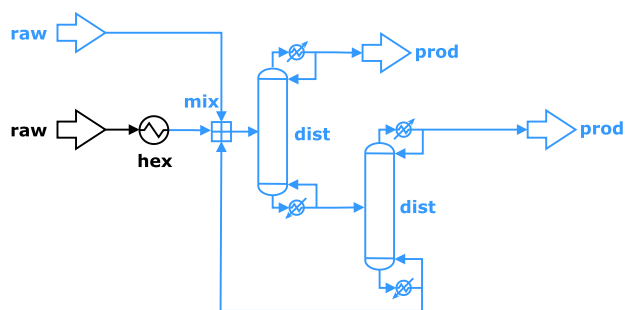
6.2. Example 2

In the second example, we consider a more complex input sequence that contains a mixer and reactor, with the outlet of the reactor being the bottom inlet of an absorption column. The goal of this example is to study if the autocompletion model can reference a recycle to the mixer in the generated tokens and, second, complete the three other streams associated with an absorption column, in specific, a top inlet, bottom outlet, and top outlet.


Start: (raw)(mix)<1(r){bin}(abs)

Beam search (descending sequence probability):


1. (raw)(mix)<1(r){bin}(abs)<&(raw){tin}&[(tout)(prod)](bout)(dist)[(tout)(prod)](bout)(dist)(tout)1(bout)(prod)



(a) 1. generated sequence



(b) 2. generated sequence



(c) 3. generated sequence

Fig. 9. Example 1: Completed flowsheets using top- p sampling.

2. (raw)(mix)<1(r){bin}(abs)<&|(raw){tin}&|[{bout}(prod)]{tout}(dist)[{bout}(prod)]{tout}(hex)1
3. (raw)(mix)<1(r){bin}(abs)<&|(raw){tin}&|[{tout}(prod)]{bout}(dist)[{tout}(prod)]{bout}(dist)[{tout}(prod)]{bout}(dist){tout}1{bout}(prod)

Top- p :

1. (raw)(mix)<1(r){bin}(abs)<&|(raw){tin}&|[{bout}(prod)]{tout}(dist)[{tout}(prod)]{bout}(pp)(hex)1
2. (raw)(mix)<1(r){bin}(abs)[{bout}(prod)]{tout}(dist)[{bout}(prod)]{tout}(mix)<&|(raw)&|(pp){tin}1
3. (raw)(mix)<1(r){bin}(abs)<&|(raw){tin}&|[{bout}(prod)]{tout}(pp)(hex){1}(dist)[{tout}(hex){1}(prod)]{bout}(hex)(comp)1

Again in this example the beam search outputs lead to grammatically correct SFILES 2.0 strings. The result demonstrates the model's ability beyond predicting the next unit operation in a flowsheet. It is able to close the recycle stream by inserting a 1 and generate all streams associated with an absorption column including a new raw material stream, as shown in Figs. 10(a) to 10(c). Similar to the previous example, beam search maximizes the final sequence probability which leads to adding only few common unit operations (only distillation columns and one heat exchanger) while completing the recycle and absorption column streams.

Flowsheet completion with top- p sampling as decoding strategy only leads to two out of three grammatically correct SFILES 2.0 strings,

i.e., examples 1 and 3, shown in Figs. 11(a) and 11(b). As expected, the generated sequences comprise unit operations that occur less often in the training data, such as pumps or compressors. Moreover, the flowsheet in Fig. 11(b) shows some logical errors. For example, a pump is applied to the outlet of the absorber and a compressor is applied to the bottom outlet of the distillation column.

6.3. Current limitations and future directions

The training results in Section 5 and the autocompletion examples in Sections 6.1 and 6.2 demonstrate for the first time, that a NLP model can learn to autocomplete chemical process topologies. In particular, the model is capable of learning typical topological patterns from the flowsheet data. However, the examples also demonstrate a number of current limitations that require future research. In example 1, a number of flowsheets are predicted based on a short given sequence of a raw material stream and a heat exchanger. While the predicted flowsheets use common patterns of unit operations, they are not yet meaningful for practical applications. Currently, the NLP model takes only topologies of flowsheets as inputs and has no information about the context (e.g., concentrations, materials, operating points). However, this information is essential to fully describe a process. Moreover, the model is not able to modify the input sequence. This flexibility is important, for instance, to perform heat integration of already used heat-exchangers or insert unit operations (e.g. mixer) at the beginning of the process. In example 2, the algorithm also suggests some sequences of unit operations that are physically wrong. For instance, a compressor is added to the bottom outlet of a distillation column.

We identify three promising directions for future research that have the potential to overcome the current limitations of our proposed approach:

1. Encode additional information about the process, such as chemical compounds, reactions, stream compositions, stream conditions (temperatures and pressures), and operating points of unit operations. This could be achieved in two ways:

- Extension of SFILES 2.0 language to encode additional information,
- Modification of the model architecture to include additional process information as numerical input features, e.g., in the head of the language model. Also our previous work shows that graph-based machine learning models can process additional flowsheet information (Stops et al., 2022).

2. Increase the quantity and quality of training data. This can be achieved in multiple ways, e.g.:

- Collect more flowsheet data including the necessary additional process information from literature and patents (Schweidtmann, 2022; Theisen et al., 2022),
- Collect and extract process information from simulation files,
- Collect data directly in a simulation software.

3. Integrate physical knowledge and rules into the machine learning approach (Venkatasubramanian, 2018; Schweidtmann et al., 2021). This could reduce the data demand and enhance prediction quality.

Finally, we suggest two explorative directions for future research. First, we suggest to extend the model architecture such that the input sequence can be adapted. This would allow to account for heat integration with previous process units. Second, we suggest to operate on a phenomena level instead of a unit operation level (Lutze et al., 2013). This could potentially facilitate learning. Also, it would allow to account for intensified unit operations. Overall, our results demonstrate that the proposed transformer model can successfully learn from

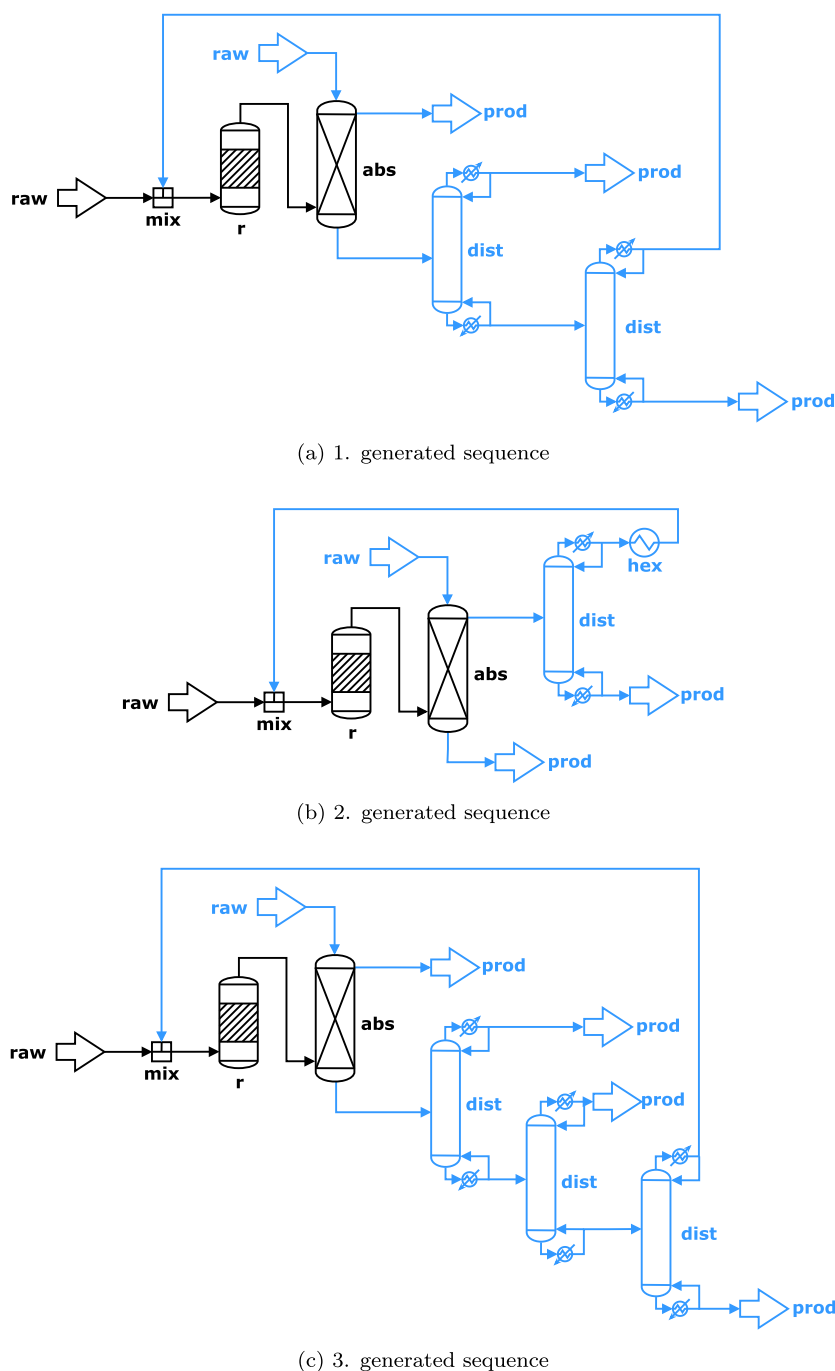


Fig. 10. Example 2: Completed flowsheets using beam search.

flowsheet topologies. Future research in the aforementioned areas has the potential to further increase the Technology Readiness Level (TRL) and pave the way for industrial applications.

7. Conclusions

We propose a novel method to learn from chemical process flowsheets and provide flowsheet structure recommendations for engineers performing process synthesis. We created two data sets, the first one consisting of synthetically generated and the second one consisting of real flowsheets in graph format. Using the conversion algorithm for the automated conversion between flowsheet graphs and SFILES 2.0 strings (Vogel et al., 2022b), we automatically generated the corresponding text-based SFILES 2.0 data sets. We trained a generative

transformer language model on over 8000 flowsheet topologies. The trained generative transformer model shows the ability to learn the grammatical structure of the SFILES 2.0 language and the patterns contained in the flowsheet topologies. Consequently, the results demonstrate that using the trained model for causal language modeling is a promising strategy to autocomplete flowsheet topologies. However, we also identified current limitations regarding applicability of the model in realistic process design scenarios.

We identify three main directions for future research: First, future work should focus on creating a larger data set of real flowsheets. To achieve this, we investigate the mining of flowsheets from literature (Schweidtmann, 2022) and the digitization of PFDs and P&IDs (Theisen et al., 2022). Also, we strive for cooperation with industrial partners to work on industrial flowsheets. Second, including

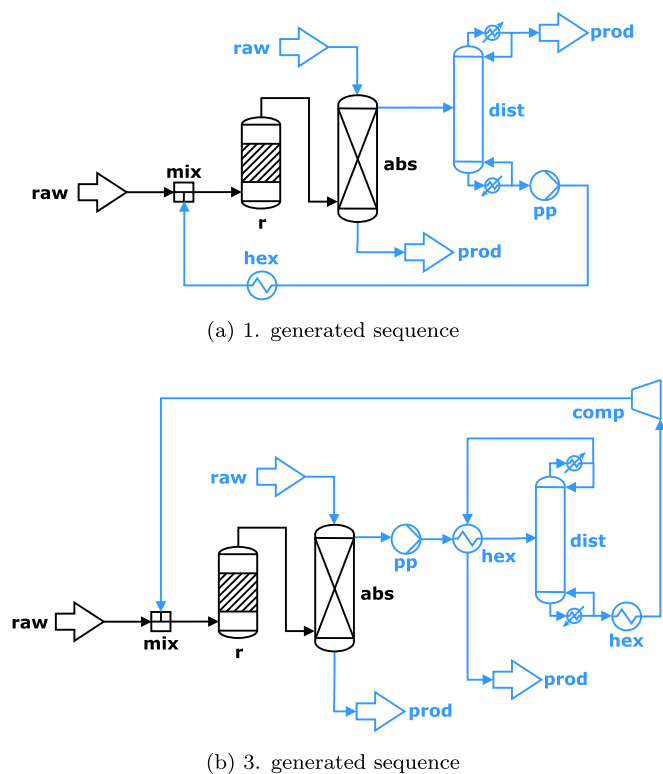


Fig. 11. Example 2: Completed flowsheets using top-*p* sampling.

the context of the chemical processes is indispensable for improving the significance and applicability of the model's predictions. This includes, for instance, chemical property data, the reaction and conversion of substances, stream compositions, and operating conditions. Third, building hybrid models by combining our data-driven approach with process design knowledge might contribute to more intelligent systems (Schweidtmann et al., 2021). Addressing these open research challenges will pave the way for the application of the proposed autocompletion method in industry. Ultimately, we envision that autocompletion of flowsheets and operating point recommendation will become a standard tool in future process simulation software.

CRedit authorship contribution statement

Gabriel Vogel: Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing – original draft, Visualization. **Lukas Schulze Balhorn:** Methodology, Validation, Formal analysis, Data curation, Writing – review & editing, Supervision. **Artur M. Schweidtmann:** Conceptualization, Methodology, Validation, Resources, Writing – review & editing, Supervision, Project administration, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The authors are unable or have chosen not to specify which data has been used.

Acknowledgments

This publication is part of the project “ChemEng KG – The Chemical Engineering Knowledge Graph” with project number 203.001.107 of the research programme “Open Science (OS) Fund 2020/2021” which is (partly) financed by the Dutch Research Council (NWO). GV acknowledges the ERASMUS Plus scholarship for his research stay at the Process Intelligence Research group. The authors acknowledge the fruitful discussions with Prof. R. Gani and Prof. J. Grievink on the combination of data-driven and mechanistic models.

References

- Biegler, L.T., Grossmann, I.E., Westerberg, A.W., 1997. *Systematic Methods of Chemical Process Design, Physical and Chemical Engineering Sciences*. Prentice Hall PTR.
- Birch, A., Osborne, M., Koehn, P., 2008. Predicting success in machine translation. In: *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pp. 745–754. URL <https://aclanthology.org/D08-1078>.
- Brown, T.B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D.M., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., Amodei, D., 2020. Language models are few-shot learners. <http://dx.doi.org/10.48550/arxiv.2005.14165>, URL <https://arxiv.org/abs/2005.14165>.
- Chen, Q., Grossmann, I.E., 2017. Recent developments and challenges in optimization-based process synthesis. *Annu. Rev. Chem. Biomol. Eng.* 8 (1), 249–283. <http://dx.doi.org/10.1146/annurev-chembioeng-080615-033546>.
- Chen, M., Tworek, J., Jun, H., Yuan, Q., d. O. Pinto, H.P., Kaplan, J., Edwards, H., Burda, Y., Joseph, N., Brockman, G., Ray, A., Puri, R., Krueger, G., Petrov, M., Khlaaf, H., Sastry, G., Mishkin, P., Chan, B., Gray, S., Ryder, N., Pavlov, M., Power, A., Kaiser, L., Bavarian, M., Winter, C., Tillet, P., Such, F.P., Cummings, D., Plappert, M., Chantzis, F., Barnes, E., Herbert-Voss, A., Guss, W.H., Nichol, A., Paino, A., Tezak, N., Tang, J., Babuschkin, I., Balaji, S., Jain, S., Saunders, W., Hesse, C., Carr, A.N., Leike, J., Achiam, J., Misra, V., Morikawa, E., Radford, A., Knight, M., Brundage, M., Murati, M., Mayer, K., Welinder, P., McGrew, B., Amodei, D., McCandlish, S., Sutskever, I., Zaremba, W., 2021. Evaluating large language models trained on code. <http://dx.doi.org/10.48550/arxiv.2107.03374>.
- Chithrananda, S., Grand, G., Ramsundar, B., 2020. Chemberta: Large-scale self-supervised pretraining for molecular property prediction. <http://dx.doi.org/10.48550/arxiv.2010.09885>, URL <https://arxiv.org/abs/2010.09885>.
- Cho, K., van Merriënboer, B., Bahdanau, D., Bengio, Y., 2014. On the properties of neural machine translation: Encoder–decoder approaches. <http://dx.doi.org/10.48550/arxiv.1409.1259>, URL <https://arxiv.org/abs/1409.1259>.
- Dai, Z., Yang, Z., Yang, Y., Carbonell, J., Le, Q.V., Salakhutdinov, R., 2019. Transformer-xl: Attentive language models beyond a fixed-length context. <http://dx.doi.org/10.48550/arxiv.1901.02860>, URL <https://arxiv.org/abs/1901.02860>.
- d'Anterrosches, L., 2006. (Ph.D. thesis). Technical University of Denmark.
- Fan, A., Lewis, M., Dauphin, Y., 2018. Hierarchical neural story generation. In: *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, <http://dx.doi.org/10.18653/v1/p18-1082>.
- Göttl, Q., Grimm, D.G., Burger, J., 2021a. Automated synthesis of steady-state continuous processes using reinforcement learning. *Front. Chem. Sci. Eng.* 16 (2), 288–302. <http://dx.doi.org/10.1007/s11705-021-2055-9>.
- Göttl, Q., Tönges, Y., Grimm, D.G., Burger, J., 2021b. Automated flowsheet synthesis using hierarchical reinforcement learning: Proof of concept. *Chem. Ing. Tech.* 93 (12), 2010–2018. <http://dx.doi.org/10.1002/cite.202100086>.
- Henao, C.A., Maravelias, C.T., 2011. Surrogate-based superstructure optimization framework. *AIChE J.* 57 (5), 1216–1232. <http://dx.doi.org/10.1002/aic.12341>.
- Hochreiter, S., Schmidhuber, J., 1997. Long short-term memory. *Neural Comput.* 9 (8), 1735–1780. <http://dx.doi.org/10.1162/neco.1997.9.8.1735>.
- Holtzman, A., Buys, J., Du, L., Forbes, M., Choi, Y., 2020. The curious case of neural text degeneration. In: *International Conference on Learning Representations*.
- Huster, W.R., Schweidtmann, A.M., Lüthje, J.T., Mitsos, A., 2020. Deterministic global superstructure-based optimization of an organic rankine cycle. *Comput. Chem. Eng.* 141, 106996. <http://dx.doi.org/10.1016/j.compchemeng.2020.106996>.
- Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., Tunyasuvunakool, K., Bates, R., Židek, A., Potapenko, A., Bridgland, A., Meyer, C., Kohl, S.A.A., Ballard, A.J., Cowie, A., Romera-Paredes, B., Nikolov, S., Jain, R., Adler, J., Back, T., Petersen, S., Reiman, D., Clancy, E., Zielinski, M., Steinegger, M., Pacholska, M., Berghammer, T., Bodensteiner, S., Silver, D., Vinyals, O., Senior, A.W., Kavukcuoglu, K., Kohli, P., Hassabis, D., 2021. Highly accurate protein structure prediction with AlphaFold. *Nature* 596 (7873), 583–589. <http://dx.doi.org/10.1038/s41586-021-03819-2>.
- Khan, A., Lapkin, A., 2020. Searching for optimal process routes: A reinforcement learning approach. *Comput. Chem. Eng.* 141, 107027. <http://dx.doi.org/10.1016/j.compchemeng.2020.107027>.

- Kim, S.H., Boukouvala, F., 2020. Surrogate-based optimization for mixed-integer nonlinear problems. *Comput. Chem. Eng.* 140, 106847. <http://dx.doi.org/10.1016/j.compchemeng.2020.106847>.
- Lakew, S.M., Cettolo, M., Federico, M., 2018. A comparison of transformer and recurrent neural networks on multilingual neural machine translation. *CoRR arXiv:1806.06957*. URL <http://arxiv.org/abs/1806.06957>.
- Lee, J.H., Shin, J., Realf, M.J., 2018. Machine learning: Overview of the recent progresses and implications for the process systems engineering field. *Comput. Chem. Eng.* 114, 111–121. <http://dx.doi.org/10.1016/j.compchemeng.2017.10.008>, FOCAP/CPC 2017, URL <https://www.sciencedirect.com/science/article/pii/S0098135417303538>.
- Lutze, P., Babi, D.K., Woodley, J.M., Gani, R., 2013. Phenomena based methodology for process synthesis incorporating process intensification. *Ind. Eng. Chem. Res.* 52 (22), 7127–7144.
- Midgley, L.I., 2020. Deep reinforcement learning for process synthesis. <http://dx.doi.org/10.48550/arxiv.2009.13265>, URL <https://arxiv.org/abs/2009.13265>.
- Oeing, J., Welscher, W., Krink, N., Jansen, L., Henke, F., Kockmann, N., 2022. Using artificial intelligence to support the drawing of piping and instrumentation diagrams using dexpi standard. *Digital Chem. Eng.* 4, 100038, URL <https://www.sciencedirect.com/science/article/pii/S2772508122000291>.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., 2018. Language models are unsupervised multitask learners. URL <https://d4mucfpxyww.cloudfront.net/better-language-models/language-models.pdf>.
- Schwaller, P., Gaudin, T., Lányi, D., Bekas, C., Laino, T., 2018. “Found in translation”: predicting outcomes of complex organic chemistry reactions using neural sequence-to-sequence models. *Chem. Sci.* 9 (28), 6091–6098. <http://dx.doi.org/10.1039/c8sc02339e>.
- Schwaller, P., Laino, T., Gaudin, T., Bolgar, P., Hunter, C.A., Bekas, C., Lee, A.A., 2019. Molecular transformer: A model for uncertainty-calibrated chemical reaction prediction. *ACS Central Sci* 5 (9), 1572–1583. <http://dx.doi.org/10.1021/acscentsci.9b00576>.
- Schwaller, P., Petraglia, R., Zullo, V., Nair, V.H., Haeuselmann, R.A., Pisoni, R., Bekas, C., Iuliano, A., Laino, T., 2020. Predicting retrosynthetic pathways using transformer-based models and a hyper-graph exploration strategy. *Chem. Sci.* 11 (12), 3316–3325. <http://dx.doi.org/10.1039/c9sc05704h>.
- Schweidtmann, A.M., 2022. Flowsheet mining. (In preparation).
- Schweidtmann, A.M., Esche, E., Fischer, A., Kloft, M., Repke, J.-U., Sager, S., Mitsos, A., 2021. Machine learning in chemical engineering: A perspective. *Chem. Ing. Tech.* 93 (12), 2029–2039. <http://dx.doi.org/10.1002/cite.202100083>.
- Schweidtmann, A.M., Mitsos, A., 2019. Deterministic global optimization with artificial neural networks embedded. *J. Optim. Theory Appl.* 180 (3), 925–948. <http://dx.doi.org/10.1007/s10957-018-1396-0>.
- Secinaro, S., Calandra, D., Secinaro, A., Muthurangu, V., Biancone, P., 2021. The role of artificial intelligence in healthcare: a structured literature review. *BMC Med. Inform. Decis. Making* 21 (1), <http://dx.doi.org/10.1186/s12911-021-01488-9>.
- Segler, M.H.S., Kogej, T., Tyrchan, C., Waller, M.P., 2017. Generating focused molecule libraries for drug discovery with recurrent neural networks. *ACS Central Sci* 4 (1), 120–131. <http://dx.doi.org/10.1021/acscentsci.7b00512>.
- Sennrich, R., Haddow, B., Birch, A., 2015. Neural machine translation of rare words with subword units. <http://dx.doi.org/10.48550/arxiv.1508.07909>, URL <https://arxiv.org/abs/1508.07909>.
- Shao, L., Gouws, S., Britz, D., Goldie, A., Strope, B., Kurzweil, R., 2017. Generating high-quality and informative conversation responses with sequence-to-sequence models. <http://dx.doi.org/10.48550/arxiv.1701.03185>, URL <https://arxiv.org/abs/1701.03185>.
- Silver, D., Huang, A., Maddison, C.J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., Hassabis, D., 2016. Mastering the game of go with deep neural networks and tree search. *Nature* 529 (7587), 484–489. <http://dx.doi.org/10.1038/nature16961>.
- Stops, L., Leenhouts, R., Gao, Q., Schweidtmann, A.M., 2022. Flowsheet synthesis through hierarchical reinforcement learning and graph neural networks. <http://dx.doi.org/10.48550/arxiv.2207.12051>, URL <https://arxiv.org/abs/2207.12051>.
- Theisen, M., Nishizaki Flores, K., Schulze Balhorn, L., Schweidtmann, A.M., 2022. Digitization of chemical process flow diagrams using deep convolutional neural networks. *Digital Chem. Eng.*
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I., 2017. Attention is all you need. URL <https://arxiv.org/pdf/1706.03762.pdf>.
- Venkatasubramanian, V., 2018. The promise of artificial intelligence in chemical engineering: Is it here finally? *AIChE J.* 65 (2), 466–478. <http://dx.doi.org/10.1002/aic.16489>.
- Vijayakumar, A.K., Cogswell, M., Selvaraju, R.R., Sun, Q., Lee, S., Crandall, D., Batra, D., 2016. Diverse beam search: Decoding diverse solutions from neural sequence models. <http://dx.doi.org/10.48550/arxiv.1610.02424>, URL <https://arxiv.org/abs/1610.02424>.
- Vogel, G., Balhorn, L.S., Hirtreiter, E., Schweidtmann, A.M., 2022a. Process-intelligence-research/sfiles2: v1.0.0. <http://dx.doi.org/10.5281/zenodo.6901932>, URL <http://dx.doi.org/10.5281/zenodo.6901932>.
- Vogel, G., Balhorn, L.S., Hirtreiter, E., Schweidtmann, A.M., 2022b. [SFILES 2.0]: An extended text-based flowsheet representation. <http://dx.doi.org/10.48550/arxiv.2208.00778>, URL <https://arxiv.org/abs/2208.00778>.
- Weber, J.M., Guo, Z., Zhang, C., Schweidtmann, A.M., Lapkin, A.A., 2021. Chemical data intelligence for sustainable chemistry. *Chem. Soc. Rev.* 50 (21), 12013–12036. <http://dx.doi.org/10.1039/d1cs00477h>.
- Weininger, D., Weininger, A., Weininger, J.L., 1989. SMILES. 2. algorithm for generation of unique SMILES notation. *J. Chem. Inf. Comput. Sci.* 29 (2), 97–101. <http://dx.doi.org/10.1021/ci00062a008>.
- Wu, Y., Schuster, M., Chen, Z., Le, Q.V., Norouzi, M., Macherey, W., Krikun, M., Cao, Y., Gao, Q., Macherey, K., Klingner, J., Shah, A., Johnson, M., Liu, X., u. Kaiser, L., Gouws, S., Kato, Y., Kudo, T., Kazawa, H., Stevens, K., Kurian, G., Patil, N., Wang, W., Young, C., Smith, J., Riesa, J., Rudnick, A., Vinyals, O., Corrado, G., Hughes, M., Dean, J., 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. <http://dx.doi.org/10.48550/arxiv.1609.08144>.
- Yang, Z., Dai, Z., Yang, Y., Carbonell, J., Salakhutdinov, R., Le, Q.V., Xlnet: Generalized autoregressive pretraining for language understanding. <http://dx.doi.org/10.48550/arxiv.1906.08237>, URL <https://arxiv.org/abs/1906.08237>.
- Zhang, T., Sahinidis, N.V., Sirola, J.J., 2018. Pattern recognition in chemical process flowsheets. *AIChE J.* 65 (2), 592–603. <http://dx.doi.org/10.1002/aic.16443>.
- Zheng, C., Chen, X., Zhang, T., Sahinidis, N.V., Sirola, J.J., 2022. Learning process patterns via multiple sequence alignment. *Comput. Chem. Eng.* 107676. <http://dx.doi.org/10.1016/j.compchemeng.2022.107676>.