Probabilistic Transformer: Modelling Ambiguities and Distributions for RNA Folding and Molecule Design

Anonymous Author(s)
Affiliation
Address
email

Abstract

Our world is ambiguous and this is reflected in the data we use to train our algorithms. This is especially true when we try to model natural processes where collected data is affected by noisy measurements and differences in measurement techniques. Sometimes, the process itself can be ambiguous, such as in the case of RNA folding, where a single nucleotide sequence can fold into multiple structures that occur with different probabilities. This ambiguity suggests that a predictive model should have similar probabilistic characteristics to match the data it models. Therefore, we propose a hierarchical latent distribution to enhance one of the most successful deep learning models, the Transformer, to accommodate these sorts of ambiguities and data distributions. We show the benefits of our approach by learning the hidden distribution on a synthetic task, with state-of-the-art results in RNA folding when training on highly ambiguous data, capable of reconstructing structure distributions and demonstrate its generative capabilities on property-based molecule design by implicitly learning the underlying property distributions and outperforming existing work.

1 Introduction

Transformer models [1] are the architecture of choice for many applications. Next to a wide range of NLP applications, such as language modelling [2, 3, 4] and machine translation [1, 5], they are also very effective in other disciplines, such as computer vision [6, 7], biology [8, 9], and chemistry [10, 11]. An additional challenging application for which transformers are promising is RNA folding, where the goal is to model a secondary structure (represented in dot-bracket notation [12]) based on a given sequence of nucleotides. RNA data is highly ambiguous since it is collected with different techniques, resolutions, protocols, and even contains natural ambiguities since there exist multiple structures for the same RNA sequence in the cell [13, 14] and the same structure can be caused by multiple sequences. Similar to RNA structures, molecules can be represented as sequences by using the *simple molecular line entry system* (SMILES) [15] and transformers arise as the architecture of choice in the field [11, 16].

Although Transformers have outstanding performance, the deterministic core of the architecture could harm performance in real-world applications like RNA folding. If collected data contains noisy labels or ambiguous samples, a vanilla Transformer model can only express uncertainties in the softmax output but not in the latent space. When sampling, sequential interdependencies can only be modelled in a decoder setting but not in an encoder-only setting. We address these limitations by proposing a probabilistic Transformer (ProbTransformer) that models a hierarchical latent distribution and performs sampling in the latent space. The ProbTransformer can represent ambiguities in these distributions and refine the sampled latent vectors within the computational graph. This is in line

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

with recent findings in cognitive science that suggest that the human brain both represents probability distributions and performs probabilistic inference [17, 18]. Our approach is based on the idea of combining the transformer architecture with a conditional variational auto-encoder (cVAE) [19], but in a hierarchical fashion similar to the hierarchical probabilistic U-Net [20]. Therefore, we introduce a new probabilistic layer and incorporate it after the attention and feed-forward layer (Section 3). In this way, we preserve the global receptive field through the attention mechanism and remain independent of other enhancements in the transformer ecosystem [21]. To train the latent distributions, we make use of the generalized ELBO with constrained optimization [22] and introduce an annealing technique to adapt a hyperparameter online.

We see our contributions in three aspects:

- The introduction of the ProbTransformer, a novel hierarchical probabilistic architecture enhancement to the Transformer ecosystem.
- Our training procedure using GECO, the analysis of the sensitivity of its hyperparameter κ, and the introduction of the online adaption technique *kappa annealing* which could be beneficial for variational training with ELBO in general.
- A comprehensive empirical analysis that verifies the ProbTransformer's capability to learn and recover data distributions on a novel synthetic sequential distribution task, assesses its capability of handling data ambiguities in practice by achieving state-of-the-art performance in RNA folding, and demonstrates its generative character by outperforming existing work in Molecule Design.

We first clarify notation and recap the cVAE [19] (Section 2), and then introduce the ProbTransformer (Section 3). We then present our experiments on the sequential distribution task (Section 4.1), RNA folding (Section 4.2), Molecule Design (Section 4.3), as well as our ablation study (Section 4.4), discuss related work (Section 5) and conclude (Section 6).

2 Background

Transformer Notation The Transformer [1] is a self-attention based sequence-to-sequence model introduced as an encoder-decoder architecture. However, both encoder-only and decoder-only versions are very successful by themselves [2, 23], and in our work, we focus on either of these two versions. An encoder or decoder has N blocks and each block consists of a multi-head (masked) attention followed by a feed-forward layer. Around both of these layers there are residual connections [24] followed by layer normalization [25]. We use the parameterization with D_{model} dimensions in the residuals and attention, H heads per attention, D_{ff} latent dimensions in the feed-forward layer, and $n \in \{1, \ldots, N\}$ for the block in the Transformer. In addition, we define $X \widehat{=} (x_1, \ldots, x_S)$ as the input sequence of length S, and $Y \widehat{=} (y_1, \ldots, y_S)$ as the target sequence, $\hat{Y} \widehat{=} (\hat{y}_1, \ldots, \hat{y}_S)$ as the predicted sequence.

Conditional Variational Auto-Encoder (cVAE) The cVAE [19] is a deep conditional generative model and an extension of the Variational Auto-Encoder [26, 27]. During inference, the cVAE aims to generate a distribution for an output $\mathbf{y}$ conditional on an input $\mathbf{x}$. More specifically, given an input $\mathbf{x}$, a latent variable $\mathbf{z}$ is drawn from a (conditional) prior model $p_\theta(\mathbf{z}|\mathbf{x})$ and used as an additional input to the generation model $p_\rho(\mathbf{y}|\mathbf{x}, \mathbf{z})$ in order to generate the prediction $\hat{\mathbf{y}}$. The cVAE is trained by minimizing the negative evidence lower bound (ELBO), $\mathcal{L}_{ELBO}$:

$$\mathcal{L}_{ELBO}(\mathbf{x}, \mathbf{y}; \theta, \rho, \psi) = \mathbb{E}_{\mathbf{z}^{post} \sim q}(-\log p_\rho(\mathbf{y}|\mathbf{x}, \mathbf{z}^{post})) + \mathrm{D}_{\mathrm{KL}}(q_\psi(\mathbf{z}|\mathbf{x}, \mathbf{y}) \parallel p_\theta(\mathbf{z}|\mathbf{x})) \quad , \quad (1)$$

where $q_\psi(\mathbf{z}|\mathbf{x}, \mathbf{y})$ describes the posterior model (called "recognition network" in [19]). All models are neural networks, and the prior and posterior models each output a mean and variance of a Gaussian distribution which represents the distribution of the latent $\mathbf{z}$. During training, $\mathbf{z}^{post} \sim q_\psi(\mathbf{z}|\mathbf{x}, \mathbf{y})$ is sampled from this posterior model and used as input to the generation model $p_\rho(\mathbf{y}|\mathbf{x}, \hat{\mathbf{z}})$, whose output $\hat{\mathbf{y}}$ is compared to the true $\mathbf{y}$ with a cross-entropy loss. This objective at training time can be viewed as a reconstruction task, which is an easier task than prediction. The KL divergence term aims to align the (conditional) prior model $p_\theta(\mathbf{z}|\mathbf{x})$ and the posterior model $q_\psi(\mathbf{z}|\mathbf{x}, \mathbf{y})$. The two losses are added and then used to train the prior, posterior, and generation models jointly, employing the reparameterization trick [28].

3 Probabilistic Transformer

Building on concepts of the cVAE, we enhance the Transformer model to a probabilistic Transformer (ProbTransformer) by adding $M \in \{1, \ldots, M\}$ probabilistic feed-forward layers (prob layer) to the existing blocks. We introduce our approach for an encoder-only model to simplify notation. It also applies to the decoder-only model but not directly to an encoder-decoder model because our training setup requires the alignment of the source and target sequences in the posterior input. Whether all or only a selection of blocks are enhanced with prob layers is a design decision and examined empirically in Section 4.4. In a block, we place the prob layer after the attention and feed-forward layer. The prob layer parameterizes a univariate Gaussian $\mathcal{N}(\mu, \sigma^2)$ used to sample a latent vector $\mathbf{z} \in \mathbb{R}^{D_z}$ for each position in the sequence. We denote by X and Y the input and target sequences of length S, by $Z_m \widehat{=} (\mathbf{z}_{m,1}, \ldots, \mathbf{z}_{m,S})$ the position-wise sampled latent sequence of block m, by $\mathbf{Z} \widehat{=} (Z_1, \ldots, Z_M)$ all latent sequences of all blocks, and by $\mathbf{Z}_{<m} \widehat{=} (Z_1, \ldots, Z_{m-1})$ all latent sequences of all blocks before m. At inference time, we sample Z_m and add it to the computation graph of the current block, similar to sampling from the (conditional) prior in the cVAE. However, in contrast to the cVAE, our sampling is conditioned hierarchically: the $\mathbf{z}$ vector also depends on previously sampled $\mathbf{z}$ vectors of all positions in all lower Transformer blocks:

$$\mathbf{z}_m \sim p(\mathbf{z}_m | \mathbf{Z}_{<m}, X). \tag{2}$$

A sequential relation is achieved by the attention mechanism: while the samples $\mathbf{z}_m$ are drawn independently for each position, the following attention operations relate them and turn the position-wise samples into a joint distribution over sequences refined in higher blocks.

As a result of this hierarchical composition of the ProbTransformer and this sequential relation, prior model $p_\theta(\mathbf{z}_m | \mathbf{Z}_{<m}, X)$ and the generation model $p_\rho(Y|X, \mathbf{Z})$ effectively become a single model $p_\Phi(Y|X)$, which strongly differs from the modeling strategy of the cVAE. This model approximately marginalizes over latent variables $\mathbf{Z}$: $P(Y|X) = \int_{\mathbf{Z}} p(\mathbf{Z}|X) \times p(Y|X, \mathbf{Z})\ d\mathbf{Z}$, and we can easily sample from it by hierarchically sampling $p(\mathbf{Z}|X)$ decomposed as

$$P_\Phi(\mathbf{Z}|X) = p(Z_M | \mathbf{Z}_{<M}, X) \cdot \ldots \cdot p(z_0|X), \tag{3}$$

followed by sampling from $p(Y|X, \mathbf{Z})$. We dub the combined model Φ the *predictive* model. At inference time, we can sample different predictions from this predictive model $\hat{Y} = \Phi(Y|X)$. However, we can also use the mean of the respective (Gaussian) distributions instead of sampling from them. We denote this as *mean inference* in contrast to *sample inference*. In the following, we introduce the prob layer architecture in detail before explaining our training setup and the learning objective.

3.1 Position-wise Probabilistic Feed-Forward Network

The prob layer is applied to each position of the latent space separately and identically. Figure 1 provides a visual description. It consists of a linear layer that transforms from the model dimension D_{model} to the distribution dimension D_z of the probabilistic latent space, followed by an activation function. Further, two linear layers with $D_z \times D_z$ weight matrices generate the latent representation of a conditional Gaussian distribution with mean $\mu_m \in \mathbb{R}^{D_z}$ and log variance[1] $\log \sigma_m^2 \in \mathbb{R}^{D_z}$ for the prob layer in block m:

$$\mu_m = \text{Linear}_{\mu,m}(\text{Act}(\text{Linear}_{In,m}(x_m))) \tag{4}$$

$$\log \sigma_m^2 = \text{Linear}_{\sigma,m}(\text{Act}(\text{Linear}_{In,m}(x_m))) \tag{5}$$

$$z_n \sim \mathcal{N}(\mu_m, \sigma_m^2) =: p(z_m | z_{<m}, X) \tag{6}$$

$$y_m = \text{Linear}_{Out,m}(z_m) \tag{7}$$

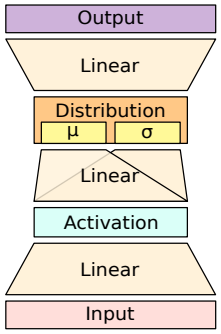

Figure 1: Probabilistic Feed-Forward layer.

As mentioned before, during *sample inference*, we can sample from the distribution (Equation 6) and during *mean inference* we use $z_n \widehat{=} \mu_m$. An additional linear layer is used to compute the layer's

[1] We use $\log \sigma_m^2$ to avoid negative variances and for numerical stability. We obtain the variance σ_m^2 as $\sigma_m^2 = e^{\log \sigma_m^2}$.

output $y_m \in \mathbb{R}^{D_{model}}$ (Equation 7). Similar to the attention and feed-forward layer, we employ a residual connection [24] followed by a layer normalization [25].

3.2 Training Setup and Learning Objective

We optimize the ELBO as the standard practice for conditional training [19]. This requires a variational posterior $Q_\psi(\mathbf{Z}|X,Y)$ that depends on both the input sequence X and target sequence Y. We model this posterior with a separate ProbTransformer Ψ with the same architecture as the predictive model except of an additional input embedding for the target sequence Y and no output generation layers, since we only make use of the latent sampling $\mathbf{Z}^{post}$ and not of its actual model output. During training, we first run the posterior model Ψ and sample in each prob layer the latent $\mathbf{Z}^{post} \sim Q_\psi$. In a second step, we run the predictive model Φ but use the latent realization Z_m^{post} instead of sampling an own Z_m in each prob layer. Figure 2 illustrates this training setup.

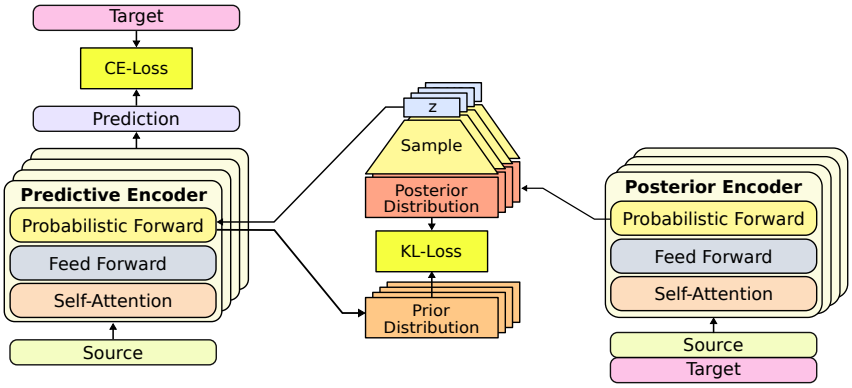

Figure 2: Training setup of the ProbTransformer: The predictive and posterior encoder trained jointly.

The negative ELBO loss $\mathcal{L}_{ELBO}$ (Equation 1) is composed of a reconstruction loss $\mathcal{L}_{rec}$ and a Kullback-Leibler divergence D_{KL} between the prior distribution P conditioned on the latent Z^{post} drawn from Q and the posterior distribution Q. The reconstruction loss is the cross-entropy between the predicted sequence $\hat{Y}$ and the true sequence Y while using the latent $\mathbf{Z}^{post}$ from the posterior model:

$$\mathcal{L}_{rec} = \mathbb{E}_{\mathbf{Z}^{post} \sim Q}\big(-\log p_\phi(\mathbf{y}|X, \mathbf{Z}^{post})\big), \tag{8}$$

and D_{KL} is the sum of KL divergences between the hierarchical decomposition, p_m and q_m, of the distributions P and Q:

$$\mathcal{D}_{KL} = \sum_{m=1}^{M} D_{\mathrm{KL}}\big(q_m(\mathbf{z}_m^{post}|\mathbf{Z}_{<m}^{post}, X, Y) \parallel p_m(\mathbf{z}_m|\mathbf{Z}_{<m}^{post}, X)\big) \tag{9}$$

With the default $\mathcal{L}_{ELBO}$ loss, we discovered instability in the training and performance issues, even with a weighting factor β [29]. Therefore, we follow [20] to avoid convergence issues with the classic ELBO objective and make use of the Generalized ELBO with Constrained Optimization (GECO) objective [22], optimizing the Lagrangian

$$\mathcal{L}_{GECO} = \lambda(\mathcal{L}_{rec} - \kappa) + \mathcal{D}_{KL}. \tag{10}$$

The Lagrange multiplier λ balances the terms and is initialized to 1, added to the learnable parameter space and updated as a function of the exponential moving average of the reconstruction loss [20]. In the beginning of the training, there is high pressure on the reconstruction loss due to an increasing λ. Once the desired reconstruction loss $\mathcal{L}_{rec} \approx \kappa$ is reached, the pressure moves to the KL-term. Please find more information about the training dynamics in Appendix A as well as more detailed illustrations of inference and training in Figure 5 and Figure 6 respectively.

κ was set to a constant in the original work of [22], but this is problematic since it is a sensitive hyperparameter due to the training dynamics. Specifically, if κ is chosen *too small*, possible failures modes include that $\mathcal{L}_{rec}$ never reaches it (with the pressure staying only on the reconstruction loss), that κ will be reached by over-fitting, or that the posterior distribution collapses which leads to high

$D_{\mathrm{KL}}(Q_\psi \parallel P_\theta)$ and destabilizes the training. On the other hand, if κ is chosen *too large*, the model can underfit or the reconstruction loss drops below κ, leading to a negative loss value and harming the training success significantly. To address this issue we introduce **kappa annealing** and adjust κ during training. Let $\mathbf{L}_c = \frac{1}{K}\sum_{k=0}^{K} \mathcal{L}_{rec}(X_k, Y_k) - \kappa$ be the mean constraint difference for K training samples in one epoch of training or a defined number of update steps. In our experiments, we find that initializing κ slightly too large and updating it every epoch with

$$\kappa = \begin{cases} \kappa + \mathbf{L}_c & \text{if } \mathbf{L}_c < 0 \text{ and } \lambda \leq 1 \\ \kappa & \text{otherwise} \end{cases} \tag{11}$$

leads to more stable training behaviour, improved final performance, and reduces the need for expensive tuning of κ.

4 Experiments

We demonstrate the benefits of the ProbTransformer in three experiments: (1) using a novel synthetic sequential distribution task, we show the advantages of distributions in the latent space over the standard softmax output distribution of a vanilla Transformer. (2) We show the benefits of the ProbTransformer in dealing with ambiguous data in RNA folding. (3) We use the ProbTransformer to improve property-based Molecule Design by sampling from the latent space instead of a softmax-output. In an additional ablation study, we provide insights on the impact of kappa annealing and hierarchical prob layers. Our source code will be publicly available upon acceptance.

4.1 Synthetic Sequential Distribution Task

In real-world generation tasks, the true distribution is not available and is only implicitly accessible by a dataset. To provide insights on the quality of distribution learning, we created a synthetic sequential distribution task and compare the ProbTransformer to two sampling methods of a vanilla Transformer.

Data We design the task to map a sequence of tokens from a source vocabulary $x \in \mathcal{V}_i^*$ to a sequence of target tokens from a target vocabulary $y \in \mathcal{V}_o^*$ with the same length. The tokens in the source sequence are used to build 'phrases" $\mathbb{P}$. Each phase exists of l tokens sampled with replacement (similar to the combination of words to phrases in a sentence). We generate randomly for each source token in each phrase a unique distribution $p(y|x, \mathbb{P})$ over the target tokens depending on the current phrase. Further, we design the distribution sparsely so that no more than k tokens from the target vocabulary have a non-zero probability. Training data are generated by sampling input sequences from all phases (with replacement) and sampling the target sequence from its corresponding distribution.

Setup We used an encoder-only model. For the ProbTransformer we enhanced each block with a prob layer. We configured the task and the model to run on one GPU within a few hours. Please find more information about the task and the configuration in Appendix B. For the inference of the vanilla Transformer we used MC dropout (using dropout during inference time, based on [30]) and sampling from the softmax distribution. For the inference of the ProbTransformer we sample from the predictive model and use the token with the highest output value. We generate 50 realizations per sample to create a predictive distribution of the target vocabulary.

Results We measure the generative performance of a model or sampling method with two metrics: The *Validity* describes the percentage of predicted tokens whose probability in the true distribution is not zero. Second, we measure the *KL-divergence* between the sampled distribution and the true distribution since sampling should reproduce the true target token distribution.As shown in Table 1, the ProbTransformer outperforms the MC dropout and softmax sampling, demonstrating its strong performance in probabilistic sequence modelling. Please note that this task does not consider interdependence sampling, where the sampling of one token in a sequence depends on the realization of others, while the next task on RNA folding does.

Table 1: The mean measures of five random seeds in the synthetic task.

Model		Validity	KL-div.
ProbTransformer		**0.99**	**0.52**
Transformer	dropout	0.93	12.71
	softmax	0.73	7.84

4.2 RNA Folding

An RNA's structure influences its function drastically [31] and the functional importance of RNA arguably is on par with that of proteins [32]. Since cellular RNAs typically have extensive secondary structures but limited tertiary structures [33], RNA folding is typically modeled as a function $\mathcal{F} : \Upsilon^* \rightarrow \Gamma^*$, where Υ and Γ denote sequence and secondary structure alphabets, respectively. However, this is a simplified view on the RNA folding process, which ignores the fact that RNAs respond to cellular conditions by altering their structures dynamically, resulting in an ensemble of structures occurring with different probabilities [13, 14] Further, real-world applications often require structure analysis of very similar sequences [34], sometimes folding into the same secondary structure. These ambiguities can hardly be captured with common approaches during training because ambiguous data is typically removed to avoid overfitting [35, 36, 37]. We address these issues with probabilistic modelling using our ProbTransformer and by keeping ambiguous training data, while implicitly measuring overfitting by explicitly removing similarity to the test data. Since we can express RNA structures in dot-bracket notation [12], the source and target have the same length, making RNA folding a suitable problem for an encoder-only Transformer model.

Data We collect a large data pool of publicly available data from recent publications [38, 39, 37, 40, 41] and split the predefined validation and test sets, VL0 and TS0 [37], from the pool. We derived a testset, *TSsameStruc*, from 149 samples of TS0 that share the same structure with at least one other sequence and uniformly sampled a disjoint set of 20 sequences with more than one annotated secondary structure from the remaining pool to produce an ambiguous testset, *TSsameSeq*. Samples without pairs and with a sequence similarity greater than 80% to the test and validation sets were removed from the pool. However, in contrast to previous work [37, 39, 42], we kept all remaining samples for training since the ProbTransformer should be strong in learning in an ambiguous setting without further data pre-processing. The final data consists of 52007 training and 1299 validation samples, and 1304 samples in the testset TS0 and 46 samples in the testset TSsameSeq. We observe 48092, 1304, and 20 unique sequences with up to 7 different structures for a single sequence, and 27179, 1204, and 46 unique structures with up to 582 different sequences for a single structure in the training, TS0, and TSsameSeq sets, respectively. This indicates that many sequences map to the same structure and different structures for the same sequence. We refer to Appendix C.1 for more information about the data.

Setup We use a 6-block encoder-only architecture with D_{model} of 512 for the vanilla Transformer and ProbTransformer. We train them for 200 epochs or 1M training steps. We compare them against the state-of-the-art deep learning algorithms, *SPOT-RNA* [37], MXFold2 [39], and a commonly-used dynamic programming approach, *RNAfold* [43] on the testset TS0 and its subset TSsameStruc using mean inference. To analyse the capabilities of the ProbTransformer to reconstruct structure distributions, we also infer the model 5, 10, 20, 50, and 100 times using sample inference on TSsameSeq, and compare against three commonly-used algorithms specialized on the prediction of structure ensembles via sampling from the Boltzmann distribution based on experimentally derived thermodynamic parameters: *RNAsubopt* [43], *RNAshapes* [44], and *RNAstructure* [45]. In contrast to previous work [37, 41, 39, 42], we do not use ensembling nor limitations on the accepted base pairs or secondary structures for post-processing, but independently train a CNN head on the output of the ProbTransformer that maps the string representation to an adjacency matrix representing the structure for evaluations on TS0 and TSsameStruc. Although we are aware of the problems with exact evaluations via F1 Score due to the dynamic nature of RNA structures [46], we report F1 Score for reasons of comparability to previous work; however, we further report the number of solved structures (the ultimate goal of the task), and the average Hamming distance to achieve a better measure of distance. For more details on the setup, CNN head, and metrics, we refer to Appendix C.2.

Results Table 2 summarizes the results of all approaches on TS0 and TSsameStruc. Both attention-based approaches generally achieve strong results, but we observe that the ProbTransformer outperforms the vanilla Transformer across all measures on both sets. The ProbTransformer achieves the best performance in terms of F1 Score, solves more than eight times the structures compared to the next best approach (MXFold2) on both sets and achieves four times lower Hamming distance on TSsameStruc compared to SPOT-RNA, indicating that it has learned to handle ambiguous sequences. Our model is further capable of accurately reconstructing challenging structures, exemplarily shown for two structures that contain long-range interactions in Figure 3. A functional important class of base pairs are so-called pseudoknots [47, 48], non-nested base pairs present in around 40% of RNAs [41] that are overrepresented in functional important regions [49, 48] and known to assist

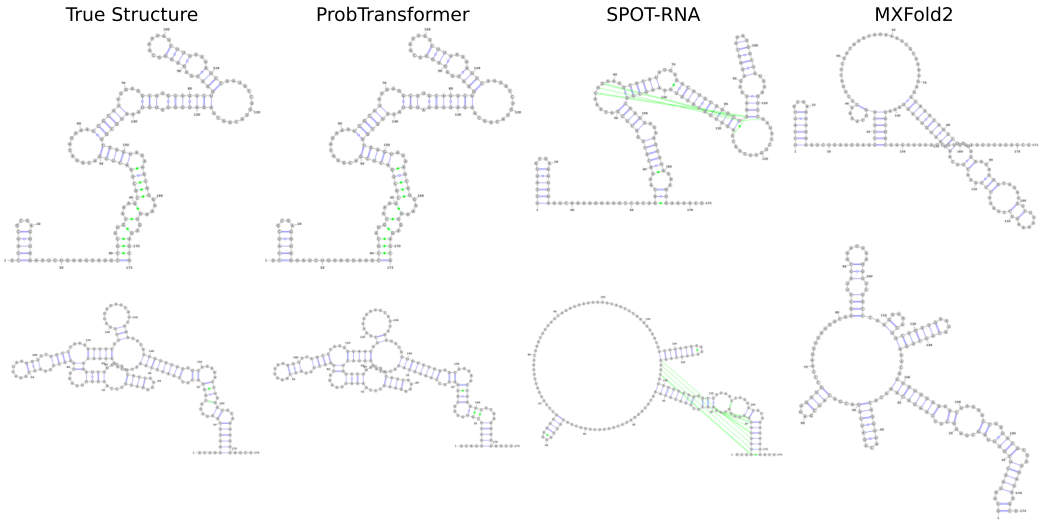

Figure 3: Example predictions of deep learning-based approaches for challenging RNAs from TS0. We show (top) a Group II catalytic intron (RF02001) and (bottom) a M-box riboswitch (RF00380).

Table 2: Structure fidelity of different RNA folding approaches on TS0 and TSsameStruc. For the ProbTransformer and the vanilla Transformer we show mean results of three random seeds.

Model	TS0			TSsameStruc		
	F1 Score	Hamming	Solved	F1 Score	Hamming	Solved
ProbTransformer	**62.5**	**27.4**	**0.118**	**93.2**	**3.2**	**0.550**
Transformer	50.5	35.3	0.084	89.5	4.6	0.481
SPOT-RNA	59.7	39.6	0.005	78.0	14.6	0.020
MXFold2	55.0	42.1	0.014	74.6	17.1	0.067
RNAFold	49.2	48.0	0.008	59.2	25.5	0.020

folding into 3D structures [50]. While RNAFold and MXFold2 cannot predict this kind of base pairs due to the underlying nearest neighbour model, Figure 3 as well as further results shown in Appendix C.3.2 indicate that the ProbTransformer can predict pseudoknots more accurately than SPOT-RNA if these are contained in the structures (in particular, see Figure 11) and further rarely predicts pseudoknots if the structure does not contain non-nested base pairs. However, a detailed analysis on the quality of pseudoknot predictions would be out of the scope of this work. When inferring the model multiple times on TSsameSeq, the ProbTransformer is the only model in our evaluation that could reconstruct more than one structure for a given sequence (for two sequences this was already achieved when inferring the model only five times). The average Hamming distances of the best predictions on TSsameSeq for every approach are summarized in Table 15 in Appendix C.3.2, showing the performance of the ProbTransformer can improve the Hamming distance by up to 44%, indicating that the ProbTransformer reconstructs the overall structure ensemble very well. Additional results and example predictions for RNA folding are reported in Appendix C.3.

4.3 Molecule Design

In the field of generative chemistry [51], deep generative models are employed to explore the chemical space [52, 53, 54]. However, biological applications typically require that the designed molecules have certain desired properties. For example, to penetrate the blood-brain barrier in order to interact with receptors of the nervous system, a molecule typically requires a certain permeability score [55, 56]. Conditional generation then refers to the problem of exploring the chemical space conditioned on molecule properties; a well-suited task to evaluate the ProbTransformer in a decoder-only setting against the state-of-the-art transformer decoder model, *MolGPT* [11].

Table 3: Multi-property (TPSA+logP+SAS) conditional training on GuacaMol dataset (mean on five different seeds).

Model	Validity	Unique	Novelty	TPSA MAD/SD	logP MAD/SD	SAS MAD/SD
ProbTransformer	**0.981**	0.821	1.0	**2.47/2.04**	**0.22/0.18**	**0.16/0.14**
MolGPT	0.973	**0.969**	1.0	3.79/4.80	0.27/0.35	0.18/0.26

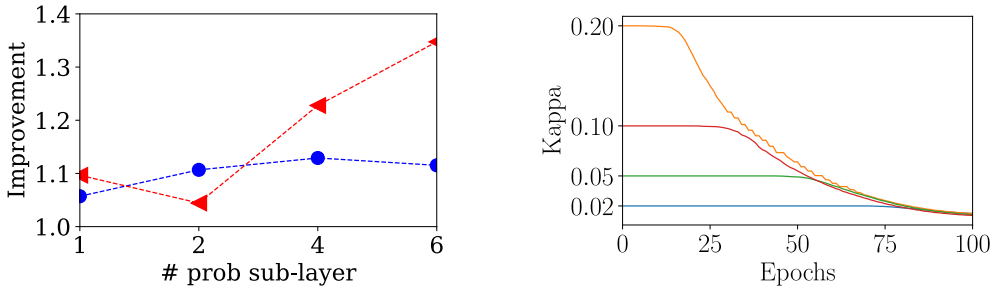

Figure 4: (Left) Performance improvement by number of prob layer: Dot (blue) on TS0 and triangle (red) on TSsameStruc. (Right) Kappa annealing with different initialization over 100 training epochs.

Data We use the GuacaMol [57] training data and the evaluation protocol provided by [11]. For more information about the data we refer to Appendix D.1.

Setup We employ the same architecture as [11], enhanced with probabilistic layers, and train our model for 20 epochs or 300k training steps. During training, the model implicitly learns the properties of the training data by conditioning generation on molecular properties together with the input SMILES. At inference, novel molecules with multiple desired property values are generated by providing the model with a start token alongside with the desired property values while predicting the next token until either a maximum length is reached or the model produces an end-token. We condition the generation of molecules on three properties: the *synthetic accessibility score (SAS)*, the *partition coefficient (logP)*, and the *topological polar surface area (TPSA)*, using the same domains of values described in [11]. For each combination of property values, the model generates a total of 10000 molecules. Results are reported in terms of the mean average deviation (MAD) and the standard deviation (SD) relative to the range of the desired property values, as well as the *validity* of the generated compounds, their *uniqueness* in terms of the internal diversity of valid predictions, and their *novelty* compared to the training data. For more information about the measures and the general setup, we refer to Appendix D.2.

Results The results for the conditional generation of molecules are summarized in Table 3. We observe a high validity score and a novelty of 1.0, indicating that our ProbTransformer has learnt the underlying SMILES grammar very well and does not suffer from overfitting to the train data. For the main task of conditional generation, the ProbTransformer clearly outperforms *MolGPT* across all measures, which highlights its ability to control multiple molecular properties during generation. The largest improvement can be observed for the SD and MAD scores of TPSA, an important measure for drug delivery in the body, improving the standard deviation (SD) by 57.5% and the mean average deviation (MAD) score by nearly 35%. We do not observe improvements in uniqueness compared to *MolGPT* but note that nearly perfect uniqueness could, e.g., be achieved by adding carbons *a posteriori* [58]. Additional results and prediction examples are shown in Appendix D.3.

4.4 Ablation Studies

We perform two ablation studies to provide more evidence for the usefulness of hierarchical probabilistic construction and kappa annealing. For this ablation, we use the RNA Folding task described in Section 4.2 with the same architecture and the training steps. We provide more details in E.

Hierarchical probabilistic design Due to the possibility of having a prob layer in any block, we tested different setups from one prob layer (middle) up to all blocks enhanced with a prob layer. In Figure 4, we measure the performance improvement in Hamming distance compared to a vanilla Transformer on the TS0 and TSsameStruct sets. Performance saturates at 4 prob layers for the TS0 set but keeps improving on the TSsameStruct set. This is in line with our hypothesis that especially ambiguous samples profit more from the hierarchical architecture.

Kappa Annealing We assess the effect of kappa annealing when initializing κ with different initial values. In this ablation, we compare a constant κ to kappa annealing on the Hamming distance, see Table 4. In these results, kappa annealing always improves the performance and substantially stabilizes final performance across a range of initialization values. Figure 4 shows the adaption of κ over a training run, demonstrating that it smoothly converges to a smaller value. Initializing with such a small value would yield much worse performance.

Table 4: Hamming distance with and w/o kappa annealing on different initialization.

Kappa	w/o	annealing
0.02	29.4	29.2
0.05	28.0	28.0
0.1	29.4	27.9
0.2	32.2	29.2

5 Related Work

Related Transformer Models A series of works enhanced the encoder-decoder Transformer architecture with a distribution in the latent space between encoder and decoder. During inference, the model samples from these distributions. There is work in the field of music representation [59] which does not use a posterior model. Work in neural machine translation [60, 61] has a separate posterior encoder and enhances the ELBO objective. This is similar to the work of Lin et al. [62] which introduces a variational decoder layer for dialogue modelling and has a separate posterior encoder. Other work in text generation [63, 64, 65] uses the encoder for the predictive and posterior model. In contrast to our work, neither of the methods mentioned above model hierarchical distributions or use the GECO objective. The work of Pei et al. [66] introduces a hierarchical stochastic multi-head attention mechanism aiming at uncertainty estimation. In contrast, we preserve the original attention mechanism by adding a new layer. The work of Liu et al. [64] also introduces a β scheduling in β-VAE ELBO objective [29] similar to our kappa annealing but focused on the KL loss instead of the reconstruction objective.

Related work in RNA folding Until recently, RNA folding was dominated by dynamic programming (DP) algorithms using either thermodynamic, statistical, or probabilistic scoring functions [67]. Learning-based approaches to the problem benefit from making few assumptions on the folding process and allowing previously unrecognized base pairing constraints [42], which recently led to state-of-the-art performance using deep learning [37, 42]. We discuss these state-of-the-art deep learning approaches in detail in Appendix C.4 and refer to a recent review [68].

Related Work in Molecule design While a plethora of deep learning-based models have been proposed for *de novo* generation of molecules in the past five years [69, 70, 71, 72, 73, 74, 75, 76, 77, 78, 79], only some methods yet approached the challenging task of generating molecules with (multiple) predefined property values (conditional generation). These include conditional RNNs [80], cVAEs [81], conditional adversarially regularized autoencoders [82], and more recently also Transformers [11]. Since [81] considers inorganic molecules and conditional generation is only performed in a case study in [82], and because of the usage of different evaluation protocols in [80] and [11], a correct evaluation and fair comparison of these approaches is out of the scope of this work. We, therefore, compare the ProbTransformer to the most similar and recent work in the field, *MolGPT*, which, to our knowledge, is the current state-of-the-art in the field. For more details on the different methods, we refer to multiple excellent reviews [83, 84, 51, 53].

6 Discussion and Conclusion

We propose a novel probabilistic layer to enhance the transformer architecture with hierarchical latent distributions while keeping the global receptive field via the attention mechanism.

We showed the benefits of our approach in several experiments. To our knowledge, the ProbTransformer is the first learned RNA folding model that can provide multiple correct structure proposals for a given RNA sequence, which opens the doors to novel research paths in RNA structure prediction that are in line with experimental evidence for RNA structural dynamics from, e.g., NMR studies, such as fraying [85], bulge migration [86], and fluctuating base pairs [86]. On the challenging multi-objective optimization task [54] of designing molecules with desired properties, we demonstrate superior control over molecule properties during the generation in a decoder-only setting compared to a state-of-the-art vanilla Transformer architecture. We want to point out that molecular research inevitably bears the risk of misuse [87], but we strongly distance ourselves from any such attempts. The ProbTransformer samples interdependent sequences in one forward path. This sampling happens in the latent space, and the ProbTransformer can refine or interpret a sampled latent representation in a further layer. Compared to sampling from the softmax output distribution, this approach yields greater flexibility. It also is compatible with other enhancements to the Transformer model since it only adds a new layer but keeps everything else unchanged.

However, our approach also has limitations. The additional layer and the posterior model increase the computational and memory needs up to a factor of two. Our approach is also limited to encoder-only or decoder-only setups since it requires a target of the same length as the input sequence. We have not applied our model in natural language processing yet, e.g. as a language model, which could improve predictive performance and save compute at inference time due to the probabilistic encoder. Another future application could be in text-to-speech or automated speech recognition since natural speech contains many ambiguities. In general, our approach and especially the conditional variational training using the GECO objective and kappa annealing, could be used in future probabilistic models. Kappa annealing reduces the need for exhaustive hyperparameter optimization and increases performance. Therefore, we expect that the positive impact of our work is not limited to RNA folding and molecule design but also to generative models in general.

References

[1] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.

[2] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.

[3] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.

[4] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*, 2022.

[5] Qiang Wang, Bei Li, Tong Xiao, Jingbo Zhu, Changliang Li, Derek F Wong, and Lidia S Chao. Learning deep transformer models for machine translation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1810–1822, 2019.

[6] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pages 8821–8831. PMLR, 2021.

[7] Salman Khan, Muzammal Naseer, Munawar Hayat, Syed Waqas Zamir, Fahad Shahbaz Khan, and Mubarak Shah. Transformers in vision: A survey. *ACM Computing Surveys (CSUR)*, 2021.

[8] Alexander Rives, Joshua Meier, Tom Sercu, Siddharth Goyal, Zeming Lin, Jason Liu, Demi Guo, Myle Ott, C Lawrence Zitnick, Jerry Ma, et al. Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences. *Proceedings of the National Academy of Sciences*, 118(15), 2021.

[9] John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589, 2021.

[10] Philippe Schwaller, Teodoro Laino, Théophile Gaudin, Peter Bolgar, Christopher A Hunter, Costas Bekas, and Alpha A Lee. Molecular transformer: a model for uncertainty-calibrated chemical reaction prediction. *ACS central science*, 5(9):1572–1583, 2019.

[11] Viraj Bagal, Rishal Aggarwal, PK Vinod, and U Deva Priyakumar. Molgpt: Molecular generation using a transformer-decoder model. *Journal of Chemical Information and Modeling*, 2021.

[12] Ivo L Hofacker, Walter Fontana, Peter F Stadler, L Sebastian Bonhoeffer, Manfred Tacker, and Peter Schuster. Fast folding and comparison of rna secondary structures. *Monatshefte für Chemie/Chemical Monthly*, 125(2):167–188, 1994.

[13] Elizabeth A Dethoff, Jeetender Chugh, Anthony M Mustoe, and Hashim M Al-Hashimi. Functional complexity and regulation through rna dynamics. *Nature*, 482(7385):322–330, 2012.

[14] Laura R Ganser, Megan L Kelly, Daniel Herschlag, and Hashim M Al-Hashimi. The roles of structural dynamics in the cellular functions of rnas. *Nature reviews Molecular cell biology*, 20(8):474–489, 2019.

[15] David Weininger. Smiles, a chemical language and information system. 1. introduction to methodology and encoding rules. *Journal of chemical information and computer sciences*, 28(1):31–36, 1988.

[16] Orion Dollar, Nisarg Joshi, David AC Beck, and Jim Pfaendtner. Attention-based generative models for de novo molecular design. *Chemical Science*, 12(24):8362–8372, 2021.

[17] Alexandre Pouget, Jeffrey M Beck, Wei Ji Ma, and Peter E Latham. Probabilistic brains: knowns and unknowns. *Nature neuroscience*, 16(9):1170–1178, 2013.

[18] Marcus Lindskog, Pär Nyström, and Gustaf Gredebäck. Can the brain build probability distributions? *Frontiers in Psychology*, 12, 2021.

[19] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. In C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc., 2015.

[20] Simon AA Kohl, Bernardino Romera-Paredes, Klaus H Maier-Hein, Danilo Jimenez Rezende, SM Ali Eslami, Pushmeet Kohli, Andrew Zisserman, and Olaf Ronneberger. A hierarchical probabilistic u-net for modeling multi-scale ambiguities. In *Medical Imaging Workshop at NeurIPS*, 2019.

[21] Yi Tay, Mostafa Dehghani, Dara Bahri, and Donald Metzler. Efficient transformers: A survey. *ACM Comput. Surv.*, apr 2022. Just Accepted.

[22] Danilo Jimenez Rezende and Fabio Viola. Taming vaes. *arXiv preprint arXiv:1810.00597*, 2018.

[23] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving language understanding by generative pre-training. *OpenAI Blog*, 2018.

[24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778. IEEE, 2016.

[25] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.

[26] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *In 2nd International Conference on Learning Representations (ICLR)*, 2013.

[27] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *International Conference on Machine Learning*, pages 1278–1286. PMLR, 2014.

[28] Diederik P Kingma, Max Welling, et al. An introduction to variational autoencoders. *Foundations and Trends® in Machine Learning*, 12(4):307–392, 2019.

[29] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. In *International Conference on Machine Learning*. PMLR, 2017.

[30] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059. PMLR, 2016.

[31] Minakshi Gandhi, Maiwen Caudron-Herger, and Sven Diederichs. Rna motifs and combinatorial prediction of interactions, stability and localization of noncoding rnas. *Nature structural & molecular biology*, 25(12):1070–1076, 2018.

[32] Sarah Djebali, Carrie A Davis, Angelika Merkel, Alex Dobin, Timo Lassmann, Ali Mortazavi, Andrea Tanzer, Julien Lagarde, Wei Lin, Felix Schlesinger, et al. Landscape of transcription in human cells. *Nature*, 489(7414):101–108, 2012.

[33] Suzanne G Rzuczek, Lesley A Colgan, Yoshio Nakai, Michael D Cameron, Denis Furling, Ryohei Yasuda, and Matthew D Disney. Precise small-molecule recognition of a toxic cug rna repeat expansion. *Nature chemical biology*, 13(2):188, 2017.

[34] Milad Miladi, Martin Raden, Sven Diederichs, and Rolf Backofen. Mutarna: analysis and visualization of mutation-induced changes in rna structure. *Nucleic acids research*, 48(W1):W287–W291, 2020.

[35] Yuedong Yang, Xiaomei Li, Huiying Zhao, Jian Zhan, Jihua Wang, and Yaoqi Zhou. Genome-scale characterization of rna tertiary structures and their functional impact by rna solvent accessibility prediction. *Rna*, 23(1):14–22, 2017.

[36] Ivantha Guruge, Ghazaleh Taherzadeh, Jian Zhan, Yaoqi Zhou, and Yuedong Yang. B-factor profile prediction for rna flexibility using support vector machines. *Journal of Computational Chemistry*, 39(8):407–411, 2018.

[37] Jaswinder Singh, Jack Hanson, Kuldip Paliwal, and Yaoqi Zhou. Rna secondary structure prediction using an ensemble of two-dimensional deep neural networks and transfer learning. *Nature communications*, 10(1):1–13, 2019.

[38] Mirela Andronescu, Vera Bereg, Holger H Hoos, and Anne Condon. Rna strand: the rna secondary structure and statistical analysis database. *BMC bioinformatics*, 9(1):1–10, 2008.

[39] Kengo Sato, Manato Akiyama, and Yasubumi Sakakibara. Rna secondary structure prediction using deep learning with thermodynamic integration. *Nature communications*, 12(1):1–9, 2021.

[40] Jaswinder Singh, Kuldip Paliwal, Tongchuan Zhang, Jaspreet Singh, Thomas Litfin, and Yaoqi Zhou. Improved rna secondary structure and tertiary base-pairing prediction using evolutionary profile, mutational coupling and two-dimensional transfer learning. *Bioinformatics*, 37, 2021.

[41] Xinshi Chen, Yu Li, Ramzan Umarov, Xin Gao, and Le Song. Rna secondary structure prediction by learning unrolled algorithms. In *International Conference on Learning Representations*, 2019.

[42] Laiyi Fu, Yingxin Cao, Jie Wu, Qinke Peng, Qing Nie, and Xiaohui Xie. Ufold: fast and accurate rna secondary structure prediction with deep learning. *Nucleic acids research*, 50(3):e14–e14, 2022.

[43] Ronny Lorenz, Stephan H. Bernhart, Christian Höner zu Siederdissen, Hakim Tafer, Christoph Flamm, Peter F. Stadler, and Ivo L. Hofacker. Viennarna package 2.0. *Algorithms for Molecular Biology*, 6(1):26, Nov 2011.

[44] Stefan Janssen, Christian Schudoma, Gerhard Steger, and Robert Giegerich. Lost in folding space? comparing four variants of the thermodynamic model for rna secondary structure prediction. *BMC bioinformatics*, 12(1):1–19, 2011.

[45] Ye Ding and Charles E Lawrence. A statistical sampling algorithm for rna secondary structure prediction. *Nucleic acids research*, 31(24):7280–7301, 2003.

[46] David H Mathews. How to benchmark rna secondary structure prediction accuracy. *Methods*, 162:60–67, 2019.

[47] Edwin Ten Dam, Kees Pleij, and David Draper. Structural and functional aspects of rna pseudoknots. *Biochemistry*, 31(47):11665–11676, 1992.

[48] David W Staple and Samuel E Butcher. Pseudoknots: Rna structures with diverse functions. *PLoS biology*, 3(6):e213, 2005.

[49] Christine E Hajdin, Stanislav Bellaousov, Wayne Huggins, Christopher W Leonard, David H Mathews, and Kevin M Weeks. Accurate shape-directed rna secondary structure modeling, including pseudoknots. *Proceedings of the National Academy of Sciences*, 110(14):5498–5503, 2013.

[50] P Fechter, J Rudinger-Thirion, C Florentz, and R Giege. Novel features in the trna-like world of plant viral rnas. *Cellular and Molecular Life Sciences CMLS*, 58(11):1547–1561, 2001.

[51] Joshua Meyers, Benedek Fabian, and Nathan Brown. De novo molecular design and generative models. *Drug Discovery Today*, 26(11):2707–2715, 2021.

[52] Christopher Lipinski and Andrew Hopkins. Navigating chemical space for biology and medicine. *Nature*, 432(7019):855–861, 2004.

[53] Ola Engkvist, Josep Arús-Pous, Esben Jannik Bjerrum, and Hongming Chen. Chapter 13 molecular de novo design through deep generative models. In *Artificial Intelligence in Drug Discovery*, pages 272–300. The Royal Society of Chemistry, 2021.

[54] Connor W Coley. Defining and exploring chemical spaces. *Trends in Chemistry*, 3(2):133–145, 2021.

[55] Han van de Waterbeemd, Gian Camenisch, Gerd Folkers, Jacques R Chretien, and Oleg A Raevsky. Estimation of blood-brain barrier crossing of drugs using molecular size and shape, and h-bonding descriptors. *Journal of drug targeting*, 6(2):151–165, 1998.

[56] David E Clark. Rapid calculation of polar molecular surface area and its application to the prediction of transport phenomena. 1. prediction of intestinal absorption. *Journal of pharmaceutical sciences*, 88(8):807–814, 1999.

[57] Nathan Brown, Marco Fiscato, Marwin HS Segler, and Alain C Vaucher. Guacamol: benchmarking models for de novo molecular design. *Journal of chemical information and modeling*, 59(3):1096–1108, 2019.

[58] Philipp Renz, Dries Van Rompaey, Jörg Kurt Wegner, Sepp Hochreiter, and Günter Klambauer. On failure modes in molecule generation and optimization. *Drug Discovery Today: Technologies*, 32:55–63, 2019.

[59] Junyan Jiang, Gus G. Xia, Dave B. Carlton, Chris N. Anderson, and Ryan H. Miyakawa. Transformer vae: A hierarchical model for structure-aware and interpretable music representation learning. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 516–520, 2020.

[60] Arya D McCarthy, Xian Li, Jiatao Gu, and Ning Dong. Improved variational neural machine translation by promoting mutual information. *arXiv preprint arXiv:1909.09237*, 2019.

[61] Arya D. McCarthy, Xian Li, Jiatao Gu, and Ning Dong. Addressing posterior collapse with mutual information for improved variational neural machine translation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8512–8525, Online, July 2020. Association for Computational Linguistics.

[62] Zhaojiang Lin, Genta Indra Winata, Peng Xu, Zihan Liu, and Pascale Fung. Variational transformers for diverse response generation. *arXiv preprint arXiv:2003.12738*, 2020.

[63] Tianming Wang and Xiaojun Wan. T-cvae: Transformer-based conditioned variational autoencoder for story completion. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 5233–5239. International Joint Conferences on Artificial Intelligence Organization, 7 2019.

[64] Danyang Liu and Gongshen Liu. A transformer-based variational autoencoder for sentence generation. In *2019 International Joint Conference on Neural Networks (IJCNN)*, pages 1–7, 2019.

[65] Le Fang, Tao Zeng, Chaochun Liu, Liefeng Bo, Wen Dong, and Changyou Chen. Transformer-based conditional variational autoencoder for controllable story generation. *arXiv preprint arXiv:2101.00828*, 2021.

[66] Jiahuan Pei, Cheng Wang, and György Szarvas. Transformer uncertainty estimation with hierarchical stochastic attention. In *Proceedings of the Thirty-Sixth AAAI Conference on Artificial Intelligence*, 2022.

[67] Elena Rivas. The four ingredients of single-sequence rna secondary structure prediction. a unifying perspective. *RNA biology*, 10(7):1185–1196, 2013.

[68] Qi Zhao, Zheng Zhao, Xiaoya Fan, Zhengwei Yuan, Qian Mao, and Yudong Yao. Review of machine learning methods for rna secondary structure prediction. *PLoS computational biology*, 17(8):e1009291, 2021.

[69] Marcus Olivecrona, Thomas Blaschke, Ola Engkvist, and Hongming Chen. Molecular de-novo design through deep reinforcement learning. *Journal of cheminformatics*, 9(1):1–14, 2017.

[70] Gabriel Lima Guimaraes, Benjamin Sanchez-Lengeling, Carlos Outeiral, Pedro Luis Cunha Farias, and Alán Aspuru-Guzik. Objective-reinforced generative adversarial networks (organ) for sequence generation models. *arXiv preprint arXiv:1705.10843*, 2017.

[71] 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). *ChemRxiv*, 2017.

[72] Artur Kadurin, Sergey Nikolenko, Kuzma Khrabrov, Alex Aliper, and Alex Zhavoronkov. drugan: an advanced generative adversarial autoencoder model for de novo generation of new molecules with desired molecular properties in silico. *Molecular pharmaceutics*, 14(9):3098–3104, 2017.

[73] Marwin HS Segler, Thierry Kogej, Christian Tyrchan, and Mark P Waller. Generating focused molecule libraries for drug discovery with recurrent neural networks. *ACS central science*, 4(1):120–131, 2018.

[74] Mariya Popova, Olexandr Isayev, and Alexander Tropsha. Deep reinforcement learning for de novo drug design. *Science advances*, 4(7):eaap7885, 2018.

[75] 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.

[76] Qi Liu, Miltiadis Allamanis, Marc Brockschmidt, and Alexander Gaunt. Constrained graph variational autoencoders for molecule design. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.

[77] Evgeny Putin, Arip Asadulaev, Quentin Vanhaelen, Yan Ivanenkov, Anastasia V Aladinskaya, Alex Aliper, and Alex Zhavoronkov. Adversarial threshold neural computer for molecular de novo design. *Molecular pharmaceutics*, 15(10):4386–4397, 2018.

[78] Oleksii Prykhodko, Simon Viet Johansson, Panagiotis-Christos Kotsias, Josep Arús-Pous, Esben Jannik Bjerrum, Ola Engkvist, and Hongming Chen. A de novo molecular generation method using latent vector based generative adversarial network. *Journal of Cheminformatics*, 11(1):1–13, 2019.

[79] Daniil Polykovskiy, Alexander Zhebrak, Benjamin Sanchez-Lengeling, Sergey Golovanov, Oktai Tatanov, Stanislav Belyaev, Rauf Kurbanov, Aleksey Artamonov, Vladimir Aladinskiy, Mark Veselov, et al. Molecular sets (moses): a benchmarking platform for molecular generation models. *Frontiers in pharmacology*, 11:1931, 2020.

[80] Panagiotis-Christos Kotsias, Josep Arús-Pous, Hongming Chen, Ola Engkvist, Christian Tyrchan, and Esben Jannik Bjerrum. Direct steering of de novo molecular generation with descriptor conditional recurrent neural networks. *Nature Machine Intelligence*, 2(5):254–265, 2020.

[81] Yashaswi Pathak, Karandeep Singh Juneja, Girish Varma, Masahiro Ehara, and U Deva Priyakumar. Deep learning enabled inorganic material generator. *Physical Chemistry Chemical Physics*, 22(46):26935–26943, 2020.

[82] Seung Hwan Hong, Seongok Ryu, Jaechang Lim, and Woo Youn Kim. Molecular generative model based on an adversarially regularized autoencoder. *Journal of chemical information and modeling*, 60(1):29–36, 2019.

[83] Benjamin Sanchez-Lengeling and Alán Aspuru-Guzik. Inverse molecular design using machine learning: Generative models for matter engineering. *Science*, 361(6400):360–365, 2018.

[84] Daniel C Elton, Zois Boukouvalas, Mark D Fuge, and Peter W Chung. Deep learning for molecular design—a review of the state of the art. *Molecular Systems Design & Engineering*, 4(4):828–849, 2019.

[85] Daniele Andreatta, Sobhan Sen, J Luis Pérez Lustres, Sergey A Kovalenko, Nikolaus P Ernsting, Catherine J Murphy, Robert S Coleman, and Mark A Berg. Ultrafast dynamics in dna:"fraying" at the end of the helix. *Journal of the American Chemical Society*, 128(21):6885–6892, 2006.

[86] Sarah A Woodson and Donald M Crothers. Proton nuclear magnetic resonance studies on bulge-containing dna oligonucleotides from a mutational hot-spot sequence. *Biochemistry*, 26(3):904–912, 1987.

[87] Fabio Urbina, Filippa Lentzos, Cédric Invernizzi, and Sean Ekins. Dual use of artificial-intelligence-powered drug discovery. *Nature Machine Intelligence*, 4(3):189–191, 2022.

[88] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.

[89] Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. Array programming with NumPy. *Nature*, 585(7825):357–362, September 2020.

[90] The pandas development team. pandas-dev/pandas: Pandas. February 2020.

[91] J. D. Hunter. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering*, 9(3):90–95, 2007.

[92] Stefan Elfwing, Eiji Uchibe, and Kenji Doya. Sigmoid-weighted linear units for neural network function approximation in reinforcement learning. *Neural Networks*, 107:3–11, 2018.

[93] John R. Prensner, Matthew K. Iyer, O. Alejandro Balbin, Saravana M. Dhanasekaran, Qi Cao, J. Chad Brenner, Bharathi Laxman, Irfan A. Asangani, Catherine S. Grasso, Hal D. Kominsky, Xuhong Cao, Xiaojun Jing, Xiaoju Wang, Javed Siddiqui, John T. Wei, Daniel Robinson, Hari K. Iyer, Nallasivam Palanisamy, Christopher A. Maher, and Arul M. Chinnaiyan. Transcriptome sequencing across a prostate cancer cohort identifies pcat-1, an unannotated lincrna implicated in disease progression. *Nature Biotechnology*, 29(8):742–749, 2011.

[94] Bingqing Cao, Tao Wang, Qiumin Qu, Tao Kang, and Qian Yang. Long noncoding rna snhg1 promotes neuroinflammation in parkinson's disease via regulating mir-7/nlrp3 pathway. *Neuroscience*, 388:118 – 127, 2018.

[95] Michael S Fernandopulle, Jennifer Lippincott-Schwartz, and Michael E Ward. Rna transport and local translation in neurodevelopmental and neurodegenerative disease. *Nature neuroscience*, 24(5):622–632, 2021.

[96] Seyedeh Hoda Alavizadeh, Maham Doagooyan, Fatemeh Zahedipour, Shima Yahoo Torghabe, Bahare Baharieh, Firooze Soleymani, and Fatemeh Gheybi. Antisense technology as a potential strategy for the treatment of coronaviruses infection: With focus on covid-19. *IET nanobiotechnology*, 2022.

[97] Lijuan Yin, Fei Zhao, Hong Sun, Zhen Wang, Yu Huang, Weijun Zhu, Fengwen Xu, Shan Mei, Xiaoman Liu, Di Zhang, et al. Crispr-cas13a inhibits hiv-1 infection. *Molecular Therapy-Nucleic Acids*, 21:147–155, 2020.

[98] Elie Dolgin et al. The tangled history of mrna vaccines. *Nature*, 597(7876):318–324, 2021.

[99] ENCODE Project Consortium et al. Identification and analysis of functional elements in 1% of the human genome by the encode pilot project. *nature*, 447(7146):799, 2007.

[100] ENCODE Project Consortium et al. An integrated encyclopedia of dna elements in the human genome. *Nature*, 489(7414):57, 2012.

[101] Ignacio Tinoco Jr and Carlos Bustamante. How rna folds. *Journal of molecular biology*, 293(2):271–281, 1999.

[102] Jörg Fallmann, Sebastian Will, Jan Engelhardt, Björn Grüning, Rolf Backofen, and Peter F Stadler. Recent advances in rna folding. *Journal of biotechnology*, 261:97–104, 2017.

[103] Padideh Danaee, Mason Rouches, Michelle Wiley, Dezhong Deng, Liang Huang, and David Hendrix. bprna: large-scale automated annotation and analysis of rna secondary structure. *Nucleic acids research*, 46(11):5381–5394, 2018.

[104] Zhen Tan, Yinghan Fu, Gaurav Sharma, and David H Mathews. Turbofold ii: Rna structural alignment and secondary structure prediction informed by multiple homologs. *Nucleic acids research*, 45(20):11570–11581, 2017.

[105] Michael F Sloma and David H Mathews. Exact calculation of loop formation probability identifies folding motifs in rna secondary structures. *RNA*, 22(12):1808–1818, 2016.

[106] Elena Rivas, Raymond Lang, and Sean R Eddy. A range of complex probabilistic models for rna secondary structure prediction that includes the nearest-neighbor model and more. *RNA*, 18(2):193–212, 2012.

[107] Limin Fu, Beifang Niu, Zhengwei Zhu, Sitao Wu, and Weizhong Li. Cd-hit: accelerated for clustering the next-generation sequencing data. *Bioinformatics*, 28(23):3150–3152, 2012.

[108] Ravi P Barnwal, Fan Yang, and Gabriele Varani. Applications of nmr to structure determination of rnas large and small. *Archives of biochemistry and biophysics*, 628:42–56, 2017.

[109] Kévin Darty, Alain Denise, and Yann Ponty. Varna: Interactive drawing and editing of the rna secondary structure. *Bioinformatics*, 25(15):1974, 2009.

[110] Jimmy Ka Ho Chiu and Yi-Ping Phoebe Chen. Efficient conversion of rna pseudoknots to knot-free structures using a graphical model. *IEEE Transactions on Biomedical Engineering*, 62(5):1265–1271, 2014.

[111] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.

[112] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.

[113] Mike Schuster and Kuldip K Paliwal. Bidirectional recurrent neural networks. *IEEE transactions on Signal Processing*, 45(11):2673–2681, 1997.

[114] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.

[115] Nathan Brown. *In Silico Medicinal Chemistry*. Theoretical and Computational Chemistry Series. The Royal Society of Chemistry, 2016.

[116] David E Clark. What has polar surface area ever done for drug discovery? *Future medicinal chemistry*, 3(4):469–484, 2011.

[117] Greg Landrum et al. Rdkit: A software suite for cheminformatics, computational chemistry, and predictive modeling, 2013.

[118] Masakazu Kondo. Developing a generative model utilizing self-attention networks: Application to materials/drug discovery. *Molecular Informatics*, 40(10):2100102, 2021.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] We describe our new approach in detail and provide a experimental evaluation with state-of-the-art results.
   (b) Did you describe the limitations of your work? [Yes] We mention them in Section 6.
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] We are not aware of negative impacts of RNA folding or foundation in probabilistic modelling but in molecule design and mention this in Section 6.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes] We read it and our paper is conform.
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] We include the source code and scripts to build the data in the supplemental material.We also open source our code upon acceptance as mentioned in Section 4
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] Details about hyperparameter and data can be found in Appendices B,C, and D.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] We report all experimental results including standard deviation on multi-seed runs in the Appendix.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] Details about compute resources can be found in Appendices B.2,C.2, and D.2.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] We cite used open source libraries in the Appendix B.2.
   (b) Did you mention the license of the assets? [No] We use only open source software.
   (c) Did you include any new assets either in the supplemental material or as a URL? [No] We do not include any new assets.
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [No]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [No] Our data does not contain personally identifiable information or offensive content.
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A] We do not use crowdsourcing nor conducted research with human subjects.
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A] We do not use crowdsourcing nor conducted research with human subjects.
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A] We do not use crowdsourcing nor conducted research with human subjects.

Appendix

A Training Dynamics of ProbTransformer

In this section, we provide insights into the training dynamics of the ProbTransformer using the RNA folding task. But first of all we recap the inference and training setup in Figure 5 and Figure 6 respectively.

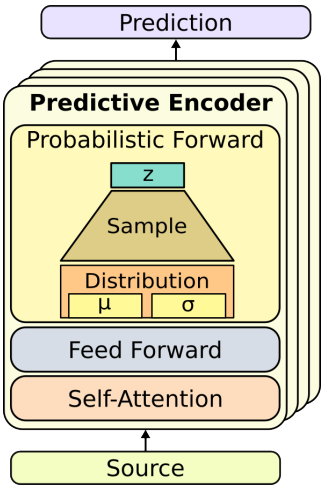

Figure 5: The figure illustrates an ProbTransformer with four blocks, each contains a Probabilistic Forward layer. This is used to model a distribution used to sample a latent z vector. The z vector is added to the computational graph via a linear layer, see Section 3.1.

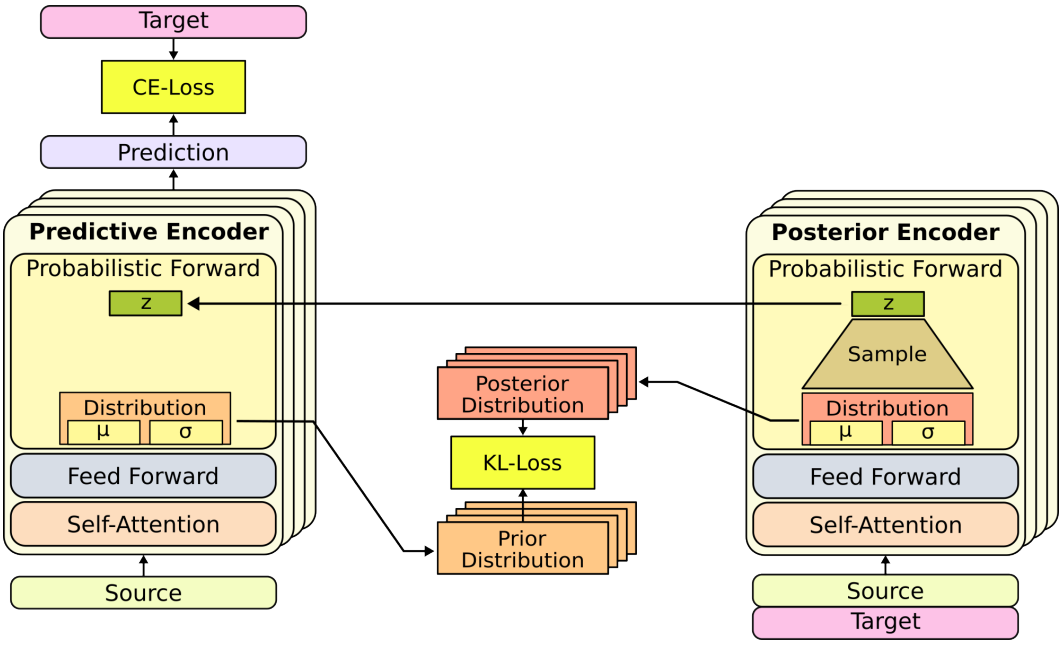

Figure 6: For the training of the ProbTransformer we use an posterior encoder with the same architecture as the predictive model but with the target sequence as additional input. During training we use the samples z vectors from the posterior encoder in the predictive encoder to generate the prediction sequence. Also, we calculate the KL-divergence between the prior distribution form the predictive encoder and the posterior distribution form the posterior encoder for each layer and sum it up. The KL-loss used in the GECO training loss to learn the distribution.

In Figure 7 we see the progress of the training. The first row pictures the hamming distance on the validation set, in the second we see the learning rate schedule, in the third the annealing of kappa, in the fourth the cross-entropy loss $\mathcal{L}_{rec}$, in the fifth the KL loss $\mathcal{D}_{KL}$, in the sixth the adjustment of κ, and at the bottom the reconstruction constraint $\mathcal{L}_{rec} - \kappa$.

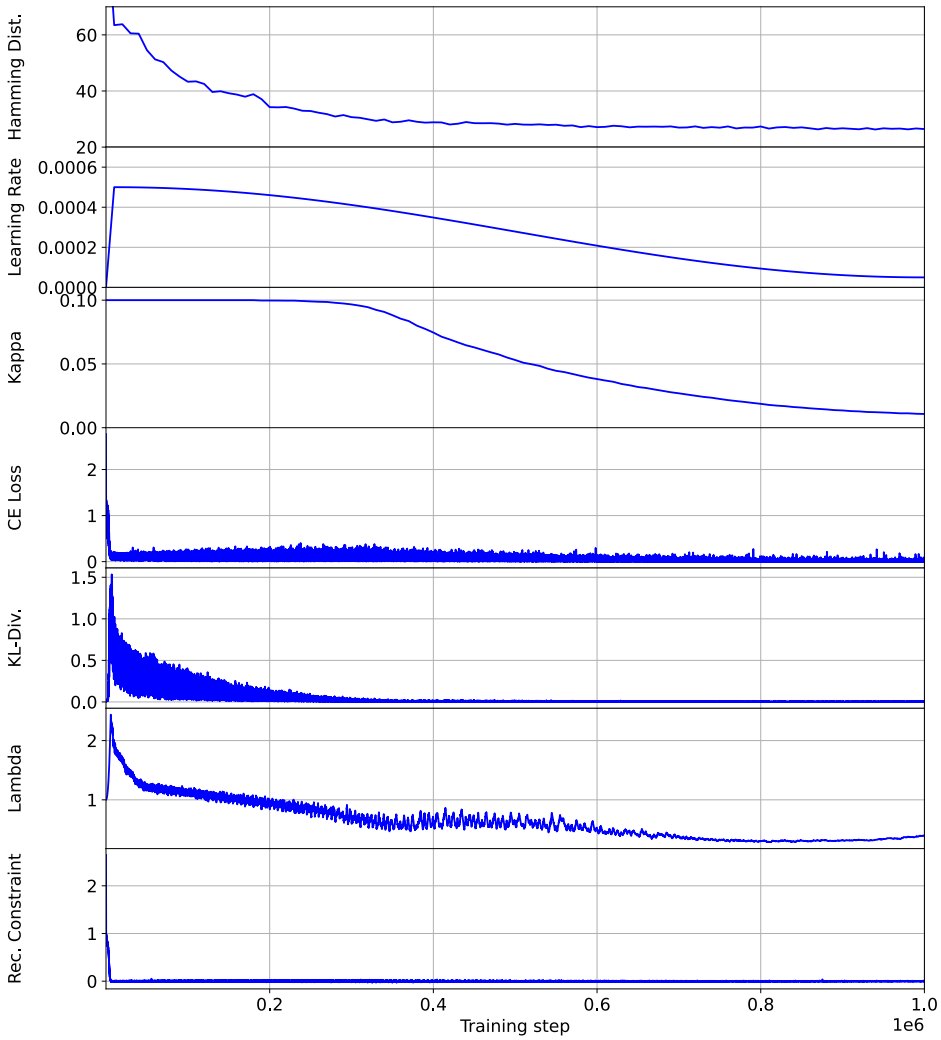

Figure 7: The training dynamics of the ProbTransformer in the RNA folding task.

We see that the κ decreases over time due to a low λ which allows to an increase in the pressure on the reconstruction. In Figure 8, we see the same training but with a log scale on the x-axis to focus on the early training phase. At the beginning of the training, the reconstruction constraint is not satisfied and the Lagrange multiplier λ is increasing which results in pressure on the reconstruction loss. At the same time, the KL divergence increases due to reconstruction via $\mathbf{Z}^{post}$, leading to an increase in the initial distance of P and Q. Once the reconstruction constraint is satisfied, λ decreases and the pressure moves to the KL term. Also, the performance, measured in terms of the Hamming distance, does not improve even when the CE loss drops, since the CE loss only trains the posterior reconstruction. It only starts to improve when the KL divergence begins to decrease and the predictive model learns to create a useful internal latent representation.

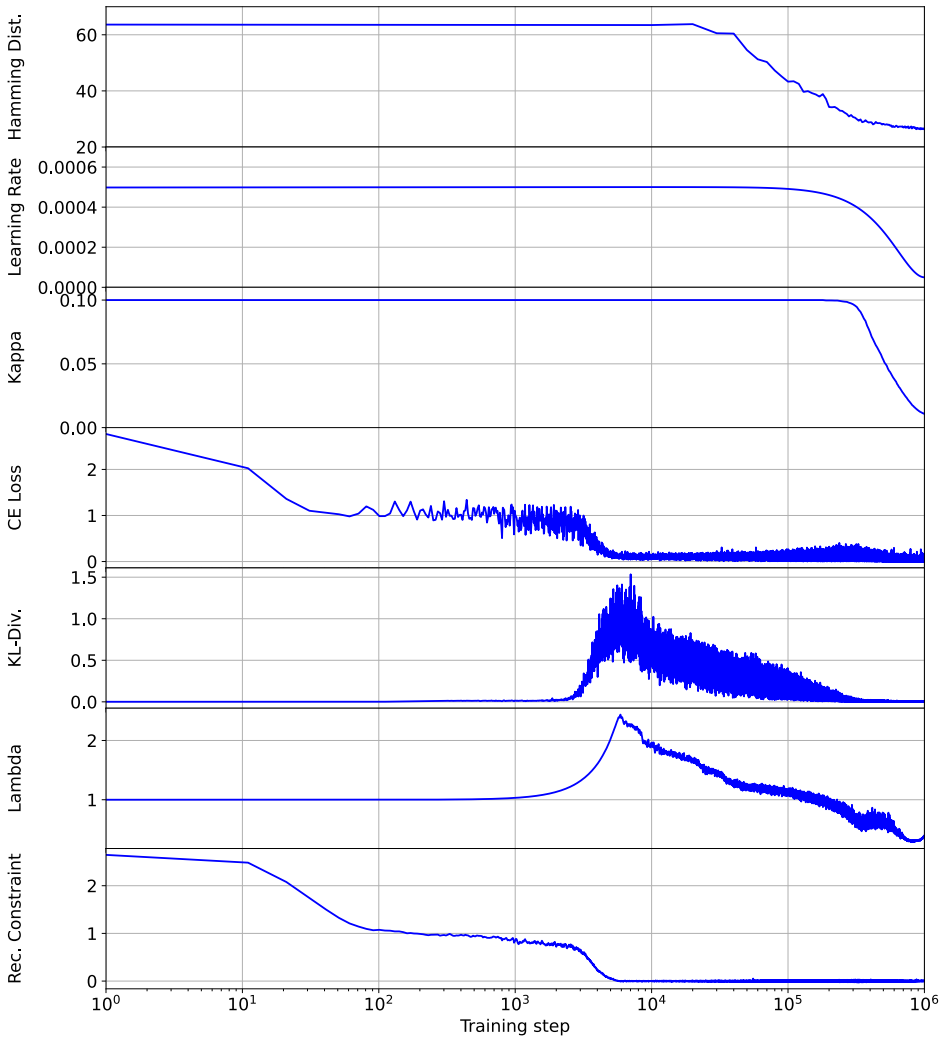

Figure 8: The training dynamics of the ProbTransformer in the RNA folding task but with log scale on the x axis to focus on the early phase of the training.

B Synthetic Sequential Distribution Task

This section provides more details about the synthetic sequential distribution task itself, the configuration of the used Transformer and ProbTransformer model, the training process, and the results.

B.1 Data

We design the task to map a sequence of tokens from a source vocabulary $x \in \mathcal{V}_i^*$ to a sequence of target tokens from a target vocabulary $y \in \mathcal{V}_o^*$ with the same length. The tokens in the source sequence are used to build 'phrases" $\mathbb{P}$. Each phrase consists of l tokens sampled with replacement (similar to the combination of words in a sentence). We randomly generate a unique distribution $p(y|x, \mathbb{P})$ over the target tokens for each source token in each phrase, depending on the current phrase. Further, we design the distribution sparsely so that no more than k tokens from the target vocabulary have a non-zero probability. The training data is generated by sampling input sequences from all phrases (with replacement) and sampling the target sequence from its corresponding distribution. The size of the source and target vocabulary is 500, a phrase exists of three tokens, and we create 1000 different sections. Each target token is drawn from a sparse distribution with 1 to 10 non-zero token probabilities. The sequence length is uniformly drawn from a length of 15 to 90. We created 100.000 training samples and 10.000 validation and test samples. Please find the detailed configuration of the task in Table 5.

Table 5: Configuration of the synthetic sequential distribution task.

Max token length	90
Min. token length	15
Number of phrases	1000
Number of training samples	100000
Token per phrase l	3
Possible target tokens k	10
Vocabulary source tokens	500
Vocabulary target tokens	500

Figure 9 shows an example of target distribution depending on the source phrase. On the x-axis, we show the target vocabulary consisting of number-tokens. On the y-axis, there are two phrases of source tokens. The yellow-green-blue color scheme represents the distribution of the target token mapping to a source token depending on the source phrase. Please note that the tokens in the second and third rows are the same but have different distributions due to the position in the phrase. A target sequence is sampled from this distribution, and in the optimal case, the model should be able to reproduce this distribution.

B.2 Setup

In general, we implement our models and tasks in Python 3.8 using mainly PyTroch [88], Numpy [89], and Pandas [90]. We use Matplotlib [91] for the plots in the paper.

For experiments on the synthetic sequential distribution task, we use the configuration listed in Table 6 for the Transformer and ProbTransformer. For MC dropout, we employed a grid search to find the optimal Dropout rate $(0.1, 0.2, \cdots, \mathbf{0.5})$. The other hyperparameters were tuned manually based on preliminary work [1] or based on preliminary experiments. We use SiLU [92] as activation function in both models. Furthermore, we use automatic mixed-precision during the training, initialize the last linear of each layer (feed-forward, attention, or prob layer) with zero, and use a learning rate warm-up in the first epoch of training as well as a cosine learning rate schedule. We use the squared softplus function to ensure a positive λ value during training and update λ with a negative gradient scaling of 0.01. For the moving average of the reconstruction loss, we use a decay of 0.95.

B.3 Results

We provide detailed results of our models and sampling methods with mean and standard deviation for five random seeds and evaluate two additional metrics: (1) We count the different output variations.

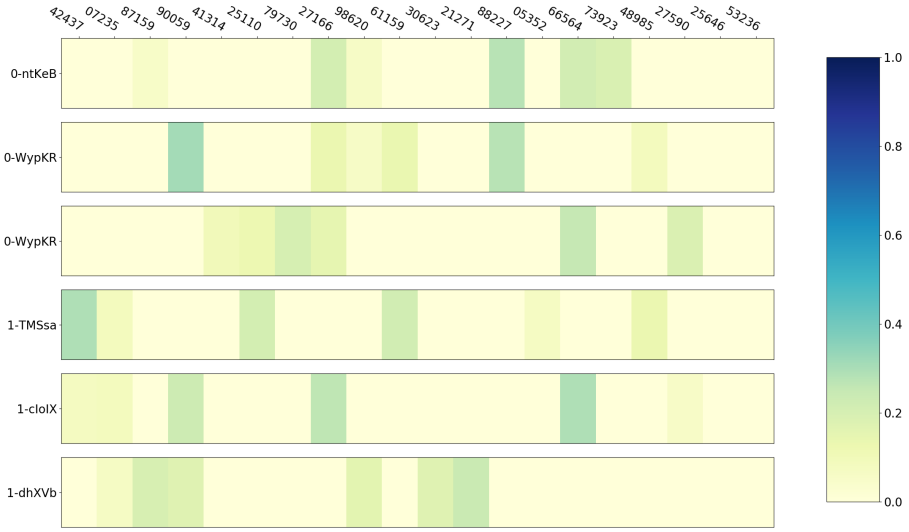

Figure 9: Target distribution depending on the source phrase.

Table 6: Hyperparameters of the Transformer and ProbTransformer training in the synthetic sequential distribution task.

Feed-forward dim	1024
Latent Z dim	256
Model dim	256
Number of layers	4
Number of heads	4
Prob layer	all layer
Kappa	0.1
Dropout	0.1
Optimizer	adamW
Beta 1	0.9
Beta 2	0.98
Gradient Clipping	100
Learning rate schedule	cosine
Learning rate high	0.001
Learning rate low	0.0001
Warmup epochs	1
Weight decay	0.01
Epochs	50
Training steps per epoch	2000

A perfect model creates the same *diversity* as nonzero probabilities in the true distribution. We normalize this measure to one; high values suggest more different tokens than non-zeros tokens in the true distribution, and smaller values suggest fewer tokens. (2) Another measure for the distance between two distributions is the *total variation* which can deal with zero probabilities. Please find the results in Table 7.

Table 7: The mean and standard deviation of five random seeds for runs with Transformer and ProbTransformer in the Synthetic Sequential Distribution Task.

Model		Validity		Diversity		KL-divergence		Total Variation	
		mean	std	mean	std	mean	std	mean	std
ProbTransformer		0.99	0.0024	0.99	0.0002	0.52	0.0165	0.11	0.0007
Transformer	dropout	0.93	0.0198	0.72	0.0014	12.71	0.0357	0.35	0.0004
	softmax	0.73	0.0075	0.90	0.0020	7.84	0.0654	0.31	0.0016

C RNA Folding

In this section, we detail our data pipeline, the general experimental setup and evaluation protocol, and show additional results, including standard deviations for multiseed runs, for our experiments on the RNA folding problem. We start, however, with a brief introduction to RNA functions and the importance of their secondary structure.

RNAs are one of the major regulators in the cell and have recently been connected to diseases like cancer [93] or Parkinson's [94, 95]. They consequently arise as a promising alternative for the development of novel drugs, including antiviral therapies against COVID-19[96] and HIV[97], or vaccines[98].

The vast majority of RNAs that are differentially transcribed from the human genome do not encode proteins [99, 100] and reveal the functions of these so-called non-coding RNAs (ncRNAs) is one of the main challenges for understanding cellular regulatory processes [31]. Similar to proteins, the function of an RNA molecule strongly depends on its folding into complex shapes, but unlike protein folding, which is dominated by hydrophobic forces acting globally, RNAs exhibit a hierarchical folding process [101]. In a first step, the corresponding nucleotides of the RNA sequence connect to each other by forming hydrogen bonds, resulting in local geometries and a distinct pairing scheme of the so-called secondary structure of RNA[2]. The secondary structure defines the accessibility of regions for interactions with other cellular compounds [31] and dictates the formation of the 3-dimensional tertiary structure [101, 102]. However, RNA structures are highly dynamic, which dramatically influences their functions [13, 14]. A learning algorithm that tackles the problem of predicting these structure ensembles is currently lacking in the field and we consider our work a major step in the direction of accurate RNA structure prediction.

C.1 Data

In this section, we detail the datasets used during training and for our experiments. RNA sequences are chains of the four nucleotides (bases) *adenine*, *cytosine*, *guanine*, and *uracil*. However, RNA data often considers an extended nucleotide alphabet using IUPAC nomenclature[3] and we note that the datasets used in this work include IUPAC nucleotides.

A RNA secondary structure is typically described as a list of pairs where a pair (i, j) denotes two nucleotides at the positions i and j of a RNA sequence that are connected by hydrogen bonds to form a base pair. In the simplest case, all pairs of the secondary structure are nested, i.e. if (i, j) and (k, l) describe two pairs of a secondary structure with $i < k$, then $i < k < l < j$. A functional important class of base pairs [47, 48], however, is called pseudoknots, where the nested pairing scheme is disrupted by one or more pairs of type: $i < k < j < l$. Canonical base pairs are formed between A and U, G and C (Watson-Crick pairs) or between G and U (Wobble pairs), while all other pairings of nucleotides are called non-canonical base pairs. We use the dot-bracket notation [12] for description of secondary structures where a dot corresponds to unpaired nucleotides and a pair of matching brackets denotes a pair of two nucleotides.

For our experiments we collect a large pool of annotated RNA secondary structures and their corresponding sequences from recent publications [41, 39, 37, 103, 38]. In particular, we collect 102098 samples from the BpRNA [103] meta database, two versions of the RNAStralign [104] dataset provided by [41] and [39] with 28168 and 20897 samples, respectively, two versions of the ArchiveII [105] dataset provided by [41] and [39] with 2936 and 3966 samples, the TR0, VL0, and TS0 datasets provided by [37] with 10814, 1300, and 1305 samples, respectively, the TrainSetA [106] and the TrainSetB [106] with 3164 and 1094 samples, respectively, and all available data from the RNA-Strand [38] database (3898 samples). For all data provided in `.bpseq`, .ct or similar file formats that only provide base pairs, we use BpRNA [103] to consistently annotate secondary structures with our major data source, the BpRNA metadatabase. We split the testset TS0 and the validation set VL0 from the pool and uniformly sampled a novel testset, TSsameSeq, from sequences of the remaining pool as described in Section 4.2. The highly redundant raw data consists of 177035 training samples,

[2]We note that there is a longstanding discussion in the field of structural biology on what is called an RNA secondary structure and we use the broadest definition of secondary structure, i.e. including non-nested structures, in this work.

[3]We refer to the IUPAC nomenclature described by the International Nucleotide Sequence Database Collaboration (INSDC) at `https://www.insdc.org/documents/feature_table.html#7.4.1`.

1300 validation samples and 1351 test samples. We remove duplicates from the data as well as samples that did not contain any pairs and resolved IUPAC nucleotides that occur very rarely in the data with a random corresponding canonical nucleotide (A, C, G, U), following the notation of the International Nucleotide Sequence Database Collaboration (INSDC) because these would unreasonably blow up the vocabulary and can hardly be learned with a deep learning approach due to their rare occurrence in the data. We note, however, that the ProbTransformer generally accepts all types of IUPAC nucleotides in the input sequence. We applied `CD-HIT-EST-2D` [107] to remove sequences from the training data with a sequence similarity greater than 80% to the validation and test samples, the lowest available threshold [37]. In accordance to [37], we limit the length of sequences to 500 nucleotides to save computational budget and since especially for longer RNAs, experimental evidence is generally still lacking because of challenges in crystallization and spectral overlap [108]. Table 8 summarizes the final datasets used for our experiments.

Table 8: Statistics of the different datasets used for our experiments on RNA folding.

Dataset	# Samples	Unique Seq.	Unique Struc.	Avg. Length	Pair Types		
					Canonical	Non-Canonical	Pseudoknots
Train	52007	48092	27179	137.46	1701469	106208	47382
VL0	1299	1299	1218	131.94	35301	4096	1001
TS0	1304	1304	1204	136.09	36702	4083	1206
TSsameStruc	149	149	49	85.04	2849	211	–
TSsameSeq	46	20	46	176.46	2273	60	150

C.2 Setup

In this section, we provide the configuration of the models and training details. We list the hyperparameters of the Transformer and ProbTransformer training in the RNA folding experiment in Table 9. We use the squared softplus function to ensure a positive λ value during training and update λ with a negative gradient scaling of 0.1. For the moving average of the reconstruction loss we use a decay of 0.95. We performed the training on one Nvidia RTX2080 GPU and the training time for one ProbTransformer model is $\sim$63h and for one Transformer $\sim$33h. The training time for the ProbTransformer nearly doubles due to the posterior model.

Table 9: Hyperparameters of the Transformer and ProbTransformer and training details of the RNA folding experiment.

Feed-forward dim	2048
Latent Z dim	512
Model dim	512
Number of layers	6
Number of heads	8
Prob layer	2,3,4,5
Kappa	0.1
Dropout	0.1
Optimizer	adamW
Beta 1	0.9
Beta 2	0.98
Gradient Clipping	100
Learning rate schedule	cosine
Learning rate high	0.0005
Learning rate low	0.00005
Warmup epochs	1
Weight decay	0.01
Epochs	100
Training steps per epoch	10000

C.2.1 CNN Head

Although the Transformer's and ProbTransformer's output prediction has a high quality, its still sometimes flawed. This hinders the evaluation of the F1 score and therefore the comparison on this metric to related work. Instead of manually designing an error correction heuristic we decided to learn a simple model which takes the Transformer's last latent and predicts an adjacency matrix which is use to evaluate the F1 score of our prediction.

We use a fixed-size CNN without up or down scaling. The detailed hyperparameters and training configuration of our CNN is listed in Table 10. The input is a concatenation of the vertical and horizontal broadcast of the last latent from the Transformer as well as the embedded nucleotide sequence. The output are two classes, one for a connection between nucleotides and one for no connection. We train our CNN head on the same training data as the Transformer and pre-compute the Transformer output to save computational resources during the CNN training, i.e. we do not train them jointly. We use early stopping based on the Hamming distance of the validation set. We perform the training on one Nvidia RTX2080 GPU and the training time is $\sim$3h.

Table 10: Hyperparameters of the CNN head and training configuration.

Model dim	64
Number of layers	8
Stride	1
Kernel	5
Dropout	0.1
Optimizer	adamW
Beta 1	0.9
Beta 2	0.98
Gradient Clipping	None
Learning rate schedule	cosine
Learning rate high	0.005
Learning rate low	0.0005
Warmup epochs	1
Weight decay	1e-10
Epochs	10
Training steps per epoch	2000

C.3 Results

In this section, we describe our evaluation protocol in detail and show additional results for the predictions on the three test sets, TS0, TSsameStruc, and TSsameSeq, including standard deviations. For drawing of RNA secondary structures, we use `VARNA` [109] provided under GNU GPL License.

C.3.1 Evaluation

We evaluate all approaches concerning Hamming distance, the number of solved tasks, and the F1 score. The Hamming distance is the raw count of mismatching characters in two strings of the same length. A dot-bracket structure with a Hamming distance of zero counts as solved. F1 score describes the harmonic mean of precision (PR) and sensitivity (SN) and is computed as follows:

$$\mathrm{PR} = \frac{\mathrm{TP}}{(\mathrm{TP} + \mathrm{FP})} \quad , \tag{12}$$

$$\mathrm{SN} = \frac{\mathrm{TP}}{(\mathrm{TP} + \mathrm{FN})} \quad , \tag{13}$$

$$\mathrm{F1} = 2 \cdot \frac{(\mathrm{PR} \cdot \mathrm{SN})}{(\mathrm{PR} + \mathrm{SN})} \quad , \tag{14}$$

where TP, FP and FN denote true positives, false positives and false negatives, respectively. For TSsameSeq, we only evaluate the best predictions concerning Hamming distance.

SPOT-RNA The output of SPOT-RNA is in `.ct` tabular format with columns for indication of pairs. Deriving pseudoknots from base pairs is not trivial [110, 103] and we, therefore, convert the output to `.bpseq` format and apply BpRNA [103], the same annotation tool as we used during data generation, to yield annotated secondary structures for all predictions of SPOT-RNA.

MXFold2 MXFold2 directly outputs secondary structures in dot-bracket format, which we evaluate directly.

RNAfold As for MXFold2, RNAfold's predictions can be evaluated directly from the output in dot-bracket format.

We note that RNAfold and MXFold2 are not capable of predicting pseudoknots due to their underlying dynamic programming approach and are limited to predictions of canonical base pairs.

UFold In contrast to all other approaches, UFold cannot handle IUPAC nucleotides in the input sequences. When evaluating the exact same test data used for all other approaches, the recommended webserver of UFold at `https://ufold.ics.uci.edu/` as well as the standalone version generates predictions with different lengths compared to the inputs which cannot be evaluated. A fair comparison with UFold thus was not possible and we decided to exclude UFold from the evaluations in the main paper. However, we resolved IUPAC nucleotides by uniformly sampling corresponding canonical nucleotides for IUPAC nucleotides to create a dataset accepted by UFold. We use the dot-bracket output of UFold for evaluations since provided `.ct` files resulted in errors when trying to obtain secondary structures using BpRNA similar to described for SPOT-RNA due to predictions with nucleotides pairing with themselves. The results of UFold on TS0 and TSsameStruc are shown in Table 11.

Table 11: Structure fidelity of UFold and the ProbTransformer on TS0 and TSsameStruc.

Model	TS0			TSsameStruc		
	F1 Score	Hamming	Solved	F1 Score	Hamming	Solved
ProbTransformer	**62.5**	**27.4**	**0.118**	**93.2**	**3.2**	**0.550**
UFold	58.8	33.3	0.038	82.7	9.4	0.141

ProbTransformer Regarding Hamming distance and the number of solved structures, we directly evaluate the predictions of the ProbTransformer from the raw model outputs. For the F1 Score, however, we use a post-processing step using the CNN head to obtain an adjacency matrix as described above.

Structure Ensemble Predictions We use the dot-bracket output of all approaches for evaluations on TSsameSeq. The predictions with the lowest Hamming distance to the respective ground truth structure were used for the evaluation of performance.

C.3.2 Detailed results

In this section, we provide further results for our experiments on the RNA folding problem, including standard derivations for multiseed runs.

TS0 We provide additional results for predictions on TS0. In Table 12 we provide results with standard deviations. Figure 10, 11 and 12 show example predictions of different approaches. Figure 13 shows the F1 score for different base-pairs in comparison with related work.

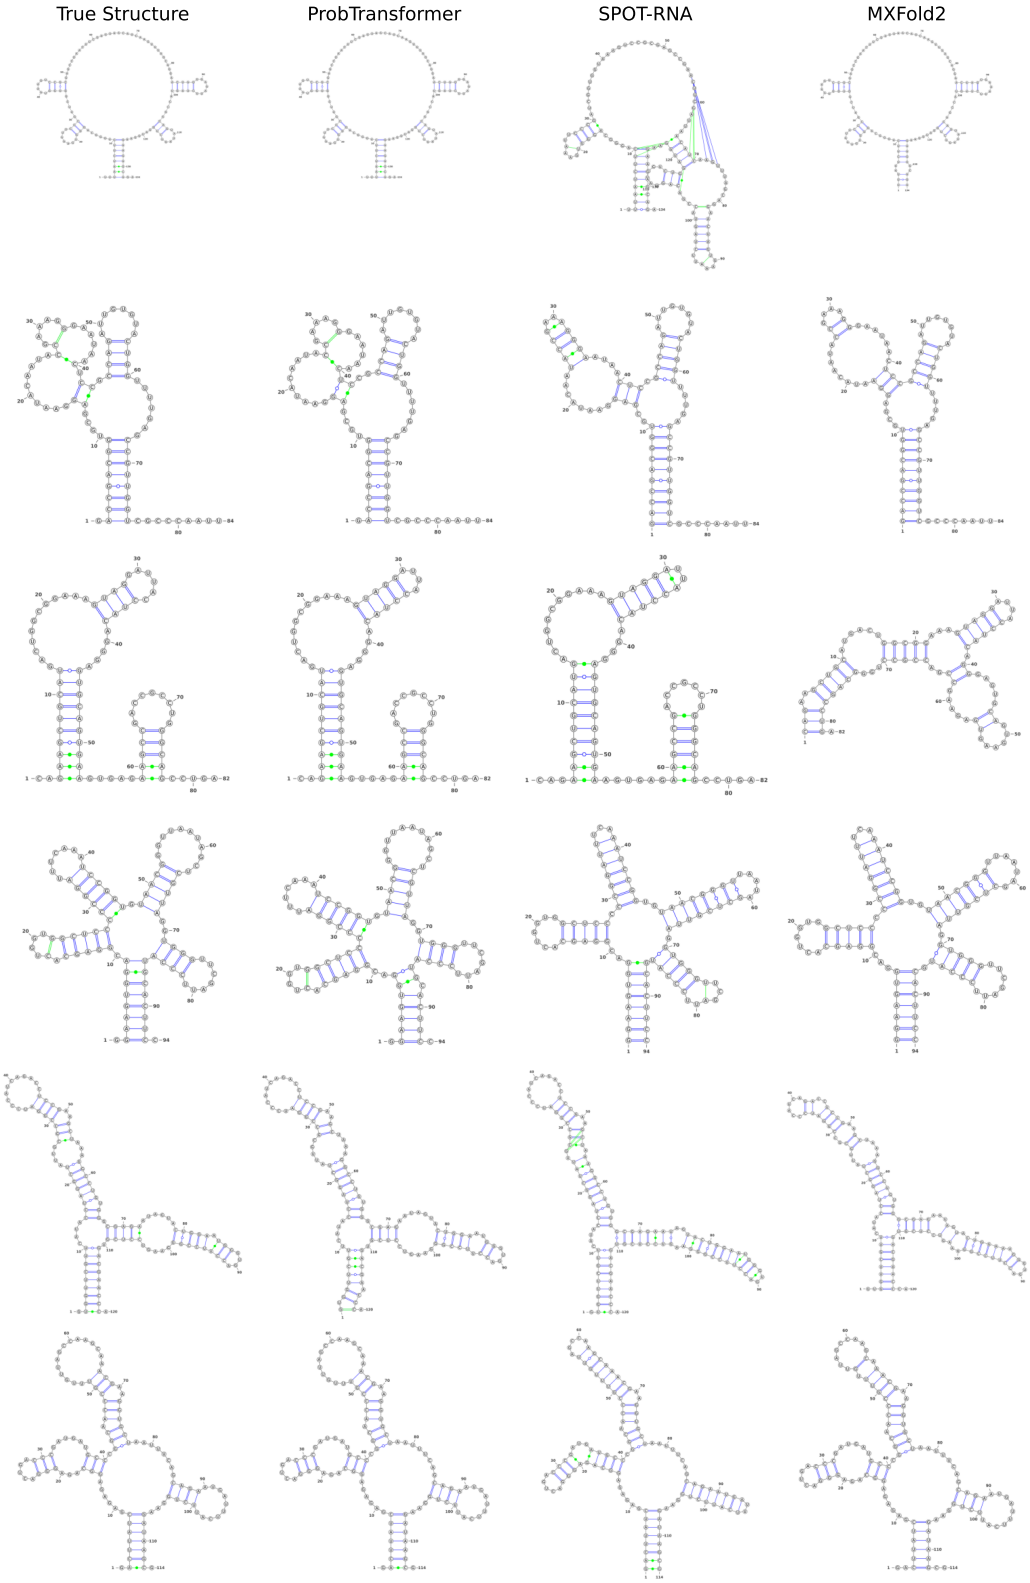

Figure 10: RNA Structure prediction examples for the test set TS0. The shown structures for the ProbTransformer are derived from the raw model outputs without further post-processing.

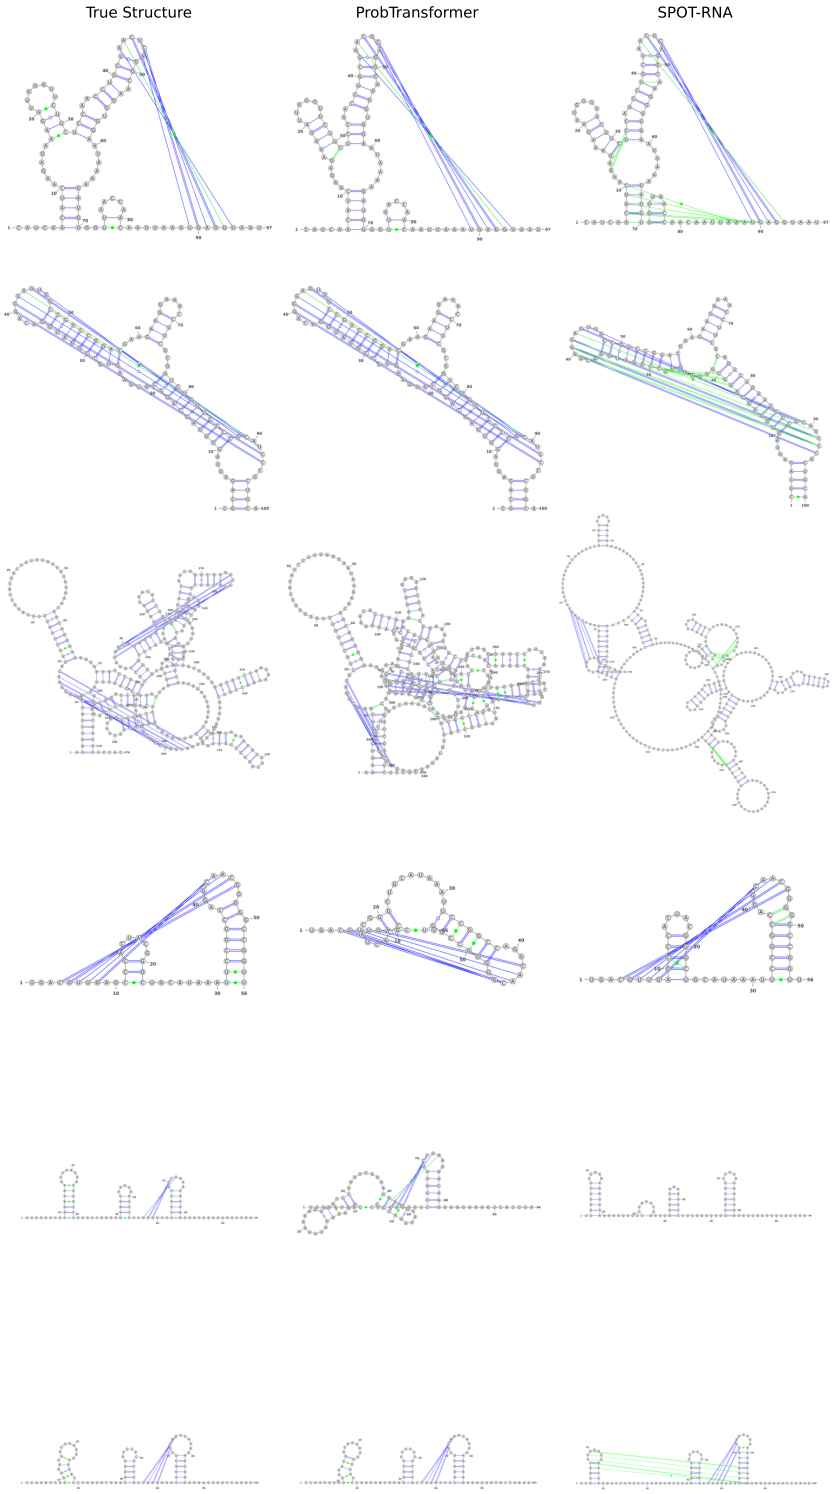

Figure 11: RNA Structure prediction examples for targets that contain pseudoknots from the test set TS0. The shown structures for the ProbTransformer are derived from the raw model outputs without further post-processing. We only show the two algorithms that are capable of predicting pseudoknots.

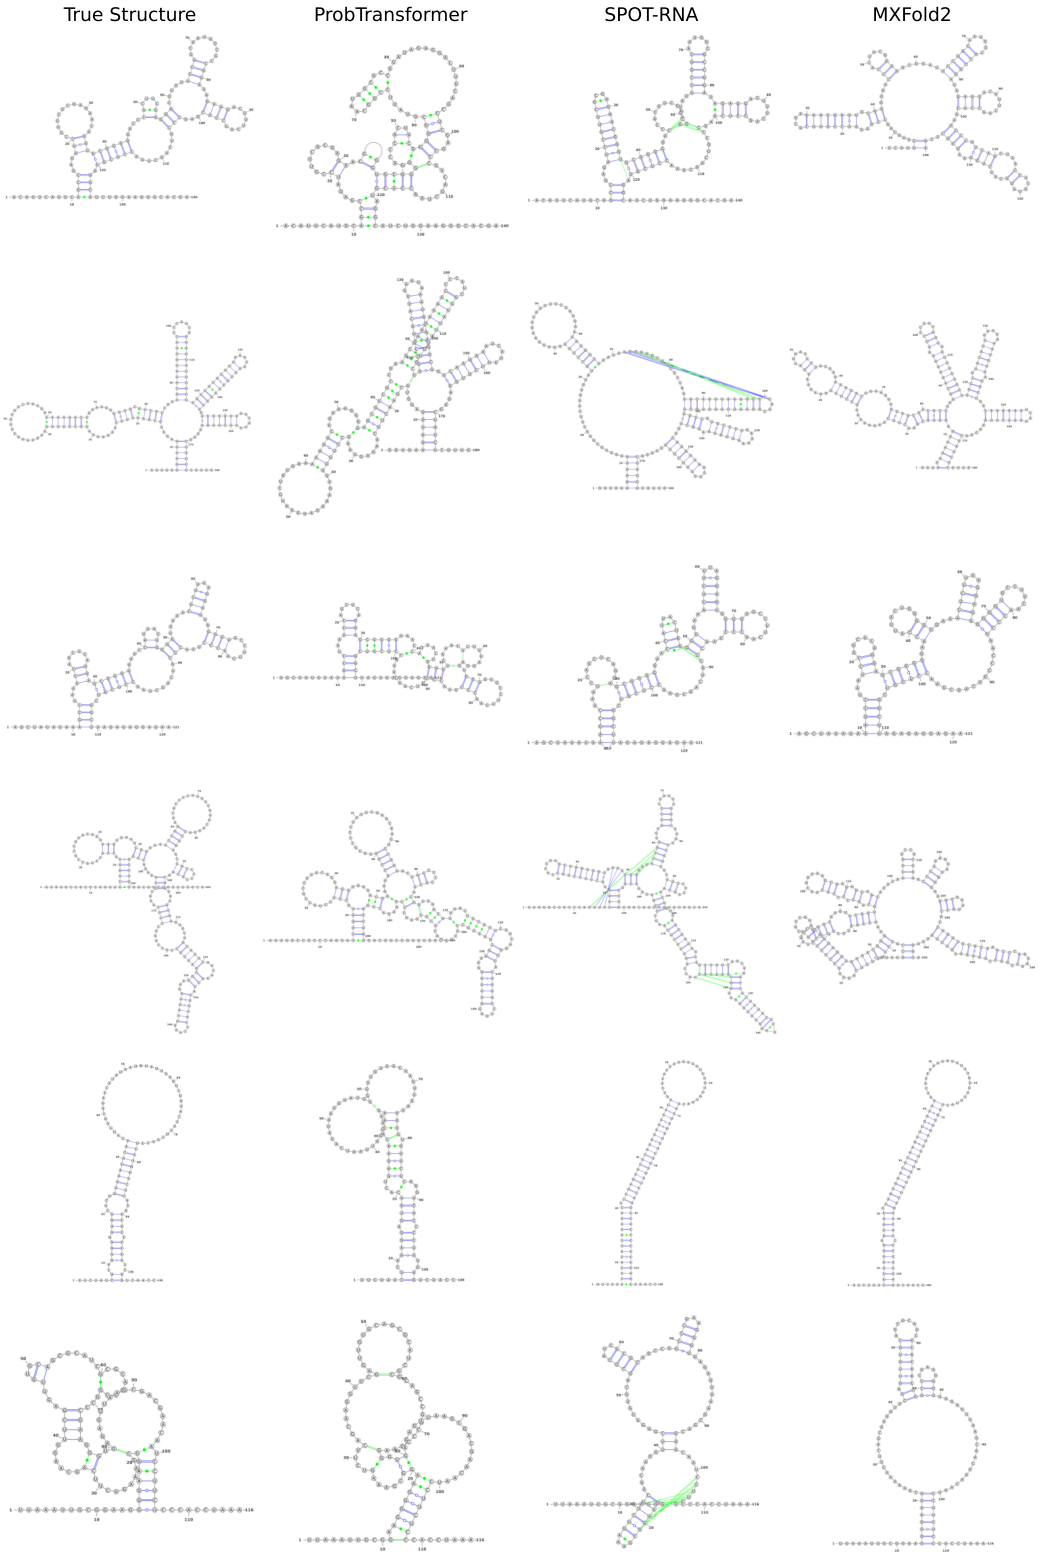

Figure 12: RNA Structure prediction examples for the test set TS0 for inaccurately predicted structures. The shown structures for the ProbTransformer are derived from the raw model outputs without further post-processing.

Table 12: Mean and standard deviation for three random seeds of the ProbTransformer and Transformer on TS0.

Model	TS0					
	F1 Score		Hamming		Solved	
	mean	std	mean	std	mean	std
ProbTransformer	62.5	0.004	27.4	0.3425	0.118	0.0037
Transformer	50.5	0.0109	35.27	0.5911	0.084	0.0019

Table 13: The F1 score of different base-pairs of the ProbTransformer, Transformer and related work on TS0.

Method	F1-All	F1-WC	F1-canonical	F1-wobble	F1-NC
ProbTransformer	62.5	65.7	64.7	58.1	39.5
Transformer	50.5	53.4	52.5	47.7	36.3
SPOT-RNA	59.7	63.7	62.7	43.9	15.6
MXFold2	55.0	58.0	57.0	41.7	0.0
RNAFold	49.2	52.0	50.8	36.9	0.0

TSsameStruc We provide additional results for predictions on TSsameStruc. Table 14 shows results with standard deviation.

Table 14: Mean and standard deviation for three random seeds of the ProbTransformer and Transformer on TSsameStruc.

Model	TSsameStruc					
	F1 Score		Hamming		Solved	
	mean	std	mean	std	mean	std
ProbTransformer	93.2	0.005	3.22	0.163	0.55	0.0055
Transformer	89.5	0.0111	4.55	0.0316	0.481	0.0084

TSsameSeq We provide results for the predictions on TSsameSeq to analyze the ability of the ProbTransformer to capture the structure distribution of RNA sequences that map to different structures. For all experiments, we inferred the model 5, 10, 20, 50, and 100 times using sample inference and analyzed the raw predictions without further post-processing. We observe that the ProbTransformer has learned the structure distributions from the data, producing predictions closer to the desired structures as indicated by a low Hamming distance shown in Table 15. The mean and standard deviations of the predictions are shown in Table 17, results for individual samples are summarized in Table 16. Remarkably, the ProbTransformer is the only model that reproduces both true structures for two of the 20 RNA sequences from the raw model predictions directly using only 5 inferences (with two out of three random seeds, results not shown).

Table 15: Average minimum Hamming distance of the different approaches on TSsameSeq.

Model	Hamming Distance				
	N=5	N=10	N=20	N=50	N=100
ProbTransformer	**26.51**	**25.16**	**24.47**	**23.60**	**23.09**
Transformer	49.17	49.17	49.17	49.17	49.17
RNAsubopt	42.59	42.83	38.09	34.22	31.30
RNAshapes	47.65	45.83	45.24	39.04	37.59
RNAstructure	47.22	42.02	38.04	35.20	32.59

Table 16: Minimal Hamming distances per Structure for all samples of TSsameSeq. We show results for one random seed of the ProbTransformer only.

Family	#Structures	ProbTransformer	RNAsubopt	RNAshapes	RNAStructure
5S rRNA	2	0/0	8/12	12/16	8/12
5S rRNA	2	8/0	16/12	14/10	24/22
5S rRNA	2	8/0	10/2	12/4	12/4
Group I catalytic intron	2	24/24	44/47	51/53	50/56
N/A	2	76/88	51/65	59/71	60/72
Antizyme RNA frameshifting stimulation element	2	14/2	12/0	16/6	15/4
N/A	3	7/5/1	6/4/0	6/4/0	6/4/0
tRNA	2	0/4	4/0	4/0	4/0
transfer-messenger RNA	3	17/33/29	91/62/112	140/111/152	107/77/119
tRNA	2	0/2	2/0	2/0	2/0
N/A	2	14/13	12/2	12/2	10/0
Bacterial small signal recognition particle RNA	2	6/7	10/6	10/6	10/6
Hammerhead ribozyme (type I)	2	36/35	14/14	8/8	8/8
Group I catalytic intron	2	118/116	54/53	102/101	66/65
5S rRNA	2	9/1	10/6	10/6	12/8
5S rRNA	2	11/1	12/12	16/8	20/14
Bacterial RNase P class A	3	14/15/15	63/60/52	55/51/46	61/59/54
Bacterial RNase P class A	3	9/6/6	50/47/42	61/58/55	43/43/38
Bacterial RNase P class B	4	65/66/68/69	78/79/101/99	80/78/104/101	65/67/89/91
5S rRNA	2	3/1	4/0	6/2	4/0

Table 17: Mean and standard deviation of the ProbTransformer and Transformer on TSsameSeq for three random seeds.

Model	TSsameSeq									
	N=5		N=10		N=20		N=50		N=100	
	mean	std	mean	std	mean	std	mean	std	mean	std
ProbTransformer	26.51	0.7754	25.16	0.4687	24.47	0.3388	23.60	0.2310	23.09	0.4820
Transformer	49.17	2.6304	49.17	2.6304	49.17	2.6304	49.17	2.6304	49.17	2.6304

C.4 Related work RNA folding

In this section, we discuss state-of-the-art deep learning approaches for the RNA folding problem in detail.

SPOT-RNA [37] was the first algorithm using deep neural networks for end-to-end prediction of RNA secondary structures. In this work, an ensemble of residual networks (ResNets) [111] and bidirectional LSTMs [112] (BiLSTMs) [113] was pre-trained on a large set of RNA secondary structure data and then fine-tuned on a small set of experimentally-derived RNA data, including tertiary interactions. Although the authors claimed the possibility of predicting RNA tertiary interactions, the performance for these types of base pairs was poor and the currently available version of the algorithm excludes tertiary interactions from its outputs. We thus consider this work as RNA secondary structure prediction.

E2efold [41] uses a Transformer encoder architecture to learn the prediction of RNA secondary structures. The algorithm was trained on a very homologous set of RNA data and showed strongly reduced performance when evaluated on data of other publications [39, 42], indicating strong overfitting. Since we use the same data set as the respective work, we exclude *E2efold* from our evaluations.

MXFold2 [39] combines deep learning with a DP approach by using a CNN/BiLSTM architecture to learn the scoring function for the DP algorithm. The network is trained to predict scores close to a set

of thermodynamic parameters to increase robustness. *MXFold2* is restricted to predict a reduced set of base pairs due to limitations in the DP algorithm.

UFold [42] employs a UNet [114] architecture for solving the RNA folding problem. Similar to *SPOT-RNA*, the authors additionally report results for predictions on data that contains tertiary interactions after fine-tuning the model on experimental data with slightly worse overall performance compared to *SPOT-RNA*. In contrast to the previously described works, however, *UFold* treats an RNA sequence as an image of all possible base-pairing maps (16 maps corresponding to 16 possible pairs) and an additional map for pair probabilities, represented as square matrices of the provided sequence.

D Molecular Design

In this section, we provide further information on our experiments for the conditional generation of molecules based on multiple desired properties.

Estimations of the size of the chemical space [52] vary widely [115] (typically between 10^{20} and 10^{100}) with a common consensus that it contains too many molecules to be explicitly enumerated [53]. Deep generative models recently attracted huge interest in exploring this practically infinite space for the use in drug discovery and deep learning-based molecular *de novo* generation has emerged as the most interesting and fast-moving field in cheminformatics [53] during the last five years. In this so-called generative chemistry [51], deep generative models are typically trained on a large part of enumerated chemical space to learn a biased distribution of molecular representations and evaluated for their ability to generate novel compounds and explore the unseen chemical space. Common evaluation protocols include metrics to measure e.g. the *novelty* of the designed compounds concerning the examples visited during training, their *uniqueness* to measure the internal diversity of predictions, and *validity* of the generated compounds regarding e.g. the underlying SMILES grammar [57]. However, besides general exploration which could be achieved using uniform sampling approaches [53], biological applications typically require that the designed molecules have certain desired properties. For generative models, the task is then to explore the chemical space conditioned on molecule properties (conditional generation).

D.1 Data

For sequence-based approaches, a common way of representing molecules is the simple molecular line-entry system (SMILES) [15]. This notation was originally proposed to represent molecules as strings and uses a sequence of elements combined with special characters to enable branching, ring-closure, and different bond orders as well as indications for properties like charges [53]. To train the ProbTransformer, we use the training data of the GuacaMol [57] benchmark suite provided by [11]. Overall, the training data consists of 1259543 SMILES with a vocabulary of 94 unique characters.

D.2 Setup

In this section, we provide the configuration of the models and training details. We list the hyperparameters of the Transformer and ProbTransformer training for the molecule design experiment in Table 18. Furthermore, we use automatic mixed-precision during the training, initialize the last linear of each layer (feed-forward, attention, or prob layer) with zero, and use a learning rate warm-up in the first epoch of training. We use the squared softplus function to ensure a positive λ value during training and update λ with a negative gradient scaling of 0.01. For the moving average of the reconstruction loss, we use a decay of 0.95. We performed the training on one Nvidia RTX2080 GPU and the training time for one ProbTransformer model is $\sim$25h and for one Transformer $\sim$13h.

We condition the generation of molecules on three properties:

The *synthetic accessibility score (SAS)* is a measure of how difficult it is to synthesize a compound. The values for *SAS* could generally range between 1 (easy to synthesize) and 10 (very difficult to make).

The *partition coefficient (logP)* describes the logarithm of the partition coefficient of a compound. This measure compares the solubilities of a solute in two immiscible solvents at equilibrium. If one of the solvents is water and the other one is non-polar, *logP* is a measure of hydrophobicity.

The *topological polar surface area (TPSA)* measures the ability of a drug to permeate cell membranes and describes the contributions of all polar atoms, such as oxygen and nitrogen and their attached hydrogens, to the molecular surface area. The polar surface area is a good estimator of the absorption, distribution, metabolism, excretion, and toxicology (ADMET)-relations of a compound and provides a rule-of-thumb for chemists to avoid dead-ends during the development process in drug discovery pipelines [116].

For our experiments, we follow the protocol described for *MolGPT* by choosing a value for each property from the following domains of values. *SAS*: 2.0, 4.0; *logP* 2.0, 6.0; *TPSA*: 40, 80. The model then conditionally generates molecules with the task to match all chosen values. Following *MolGPT*,

Table 18: Hyperparameters of the Transformer and ProbTransformer and training details of the molecule design experiment.

Feed-forward dim	1024
Latent Z dim	64
Model dim	256
Number of layers	8
Number of heads	8
Prob layer	2-7
Kappa	0.1
Dropout	0.1
Optimizer	adamW
Beta 1	0.9
Beta 2	0.98
Gradient Clipping	100
Learning rate schedule	cosine
Learning rate high	0.0005
Learning rate low	0.00005
Warmup epochs	1
Weight decay	0.01
Epochs	60
Training steps per epoch	5000

results are reported in terms of the mean average deviation (MAD) and the standard deviation (SD) relative to the range of the desired property values. While our experiment focuses on conditional generation to match the desired property values, we also report the following scores.

Validity describes the fraction of generated molecules that are expressed as valid SMILES. High validity indicates strong learning of the underlying SMILES grammar.

Uniqueness is a measure of the prediction diversity. Uniqueness is the fraction of unique predictions from all generated valid SMILES. Although this metric is known to be not well-defined [16] and can be tricked by very simple means [58], we report uniqueness scores for reasons of comparability to previous work in the field.

Novelty is the fraction of valid molecules that are different from the training samples. A high novelty indicates strong exploration while a low novelty indicates overfitting.

We use `rdkit` [117] for computations of TPSA and logP and use the provided script by [11] for computations of SAS.

D.3 Results

We provide detailed results with standard derivation in Table 19. Figure 13 shows the distributions of properties for all valid predictions of the ProbTransformer and MolGPT. We observe less deviation from the desired property values for the predictions of the ProbTransformer compared to MolGPT. In line with these results, we observe that the ProbTransformer generates more unique molecules with properties close to the desired property values compared to MolGPT as indicated in Table 20. Example predictions of these molecules are shown in Figures 14, 15, 16, 17, 18, 19, 20, and 21. We use `rdkit` [117] for drawing of the molecules.

Table 19: Results of the ProbTransformer in multi-property (TPSA+logP+SAS) conditional training on GuacaMol dataset on five different seeds.

				TPSA		logP		SAS	
	Validity	Unique	Novelty	MAD	SD	MAD	SD	MAD	SD
mean	0.981	0.821	1	2.47	2.04	0.22	0.18	0.16	0.14
std	0.0123	0.0858	0	0.3727	0.3715	0.0544	0.0449	0.0767	0.0623

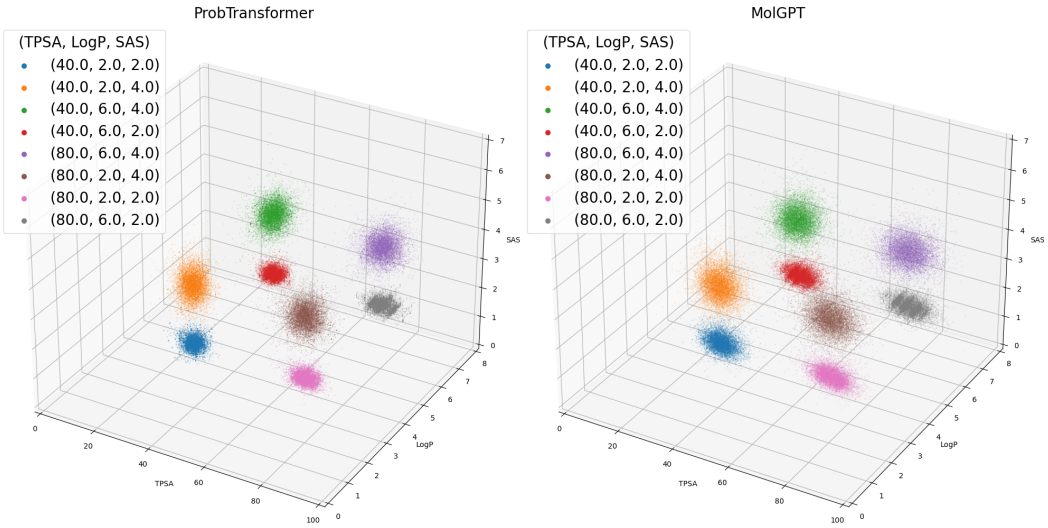

Figure 13: Distributions of predicted properties of the ProbTransformer and MolGPT. Results are shown for a representative seed for the ProbTransformer.

Table 20: Number of unique molecules that meet the desired properties for the ProbTransformer and MolGPT. We allow a deviation from the desired values of 0.5 for TPSA and 0.1 for SAS and LogP. Results show the mean with standard deviation for five random seeds of the ProbTransformer.

(TPSA, LogP, SAS)	MolGPT	ProbTransformer	
		Mean	Std
(40.0, 2.0, 2.0)	58.0	**80.4**	4.1280
(40.0, 2.0, 4.0)	25.0	**48.6**	4.8415
(40.0, 6.0, 4.0)	20.0	**40.4**	3.7202
(40.0, 6.0, 2.0)	44.0	**102.4**	9.7693
(80.0, 2.0, 2.0)	105.0	**220.2**	16.8333
(80.0, 2.0, 4.0)	49.0	**55.4**	4.4091
(80.0, 6.0, 4.0)	50.0	**60.8**	6.2418
(80.0, 6.0, 2.0)	246.0	**430.8**	57.9876

TPSA: 40.0, LogP: 2.0, SAS: 2.0

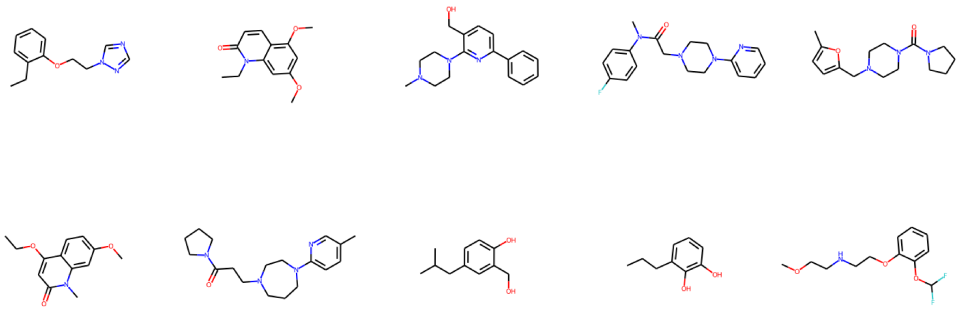

Figure 14: Example predictions of the ProbTransformer for property values of TPSA: 40.0; LogP: 2.0; SAS: 2.0 with an allowed deviation of 0.5, 0.1, and 0.1, respectively.

TPSA: 40.0, LogP: 2.0, SAS: 4.0

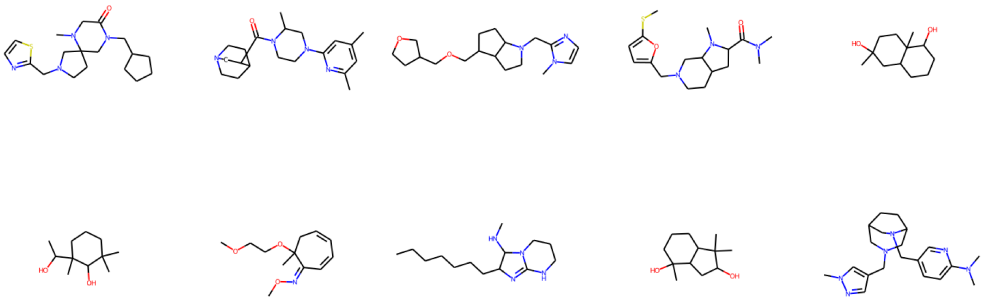

Figure 15: Example predictions of the ProbTransformer for property values of TPSA: 40.0; LogP: 2.0; SAS: 4.0 with an allowed deviation of 0.5, 0.1, and 0.1, respectively.

TPSA: 40.0, LogP: 6.0, SAS: 4.0

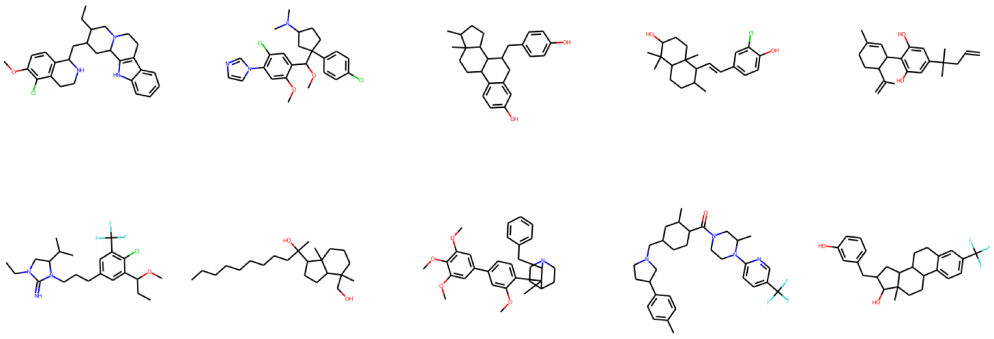

Figure 16: Example predictions of the ProbTransformer for property values of TPSA: 40.0; LogP: 6.0; SAS: 4.0 with an allowed deviation of 0.5, 0.1, and 0.1, respectively.

TPSA: 40.0, LogP: 6.0, SAS: 2.0

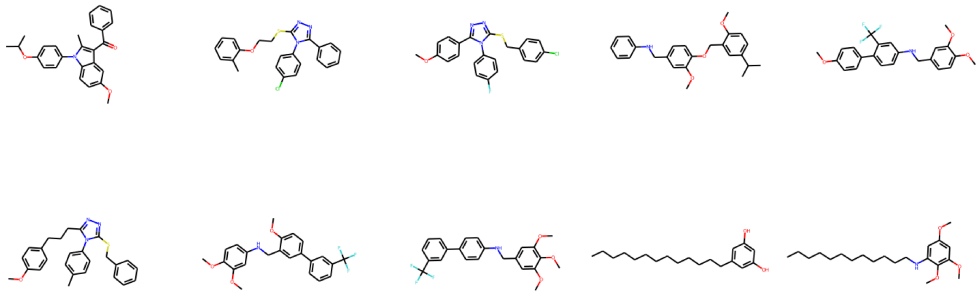

Figure 17: Example predictions of the ProbTransformer for property values of TPSA: 40.0; LogP: 6.0; SAS: 2.0 with an allowed deviation of 0.5, 0.1, and 0.1, respectively.

TPSA: 80.0, LogP: 2.0, SAS: 2.0

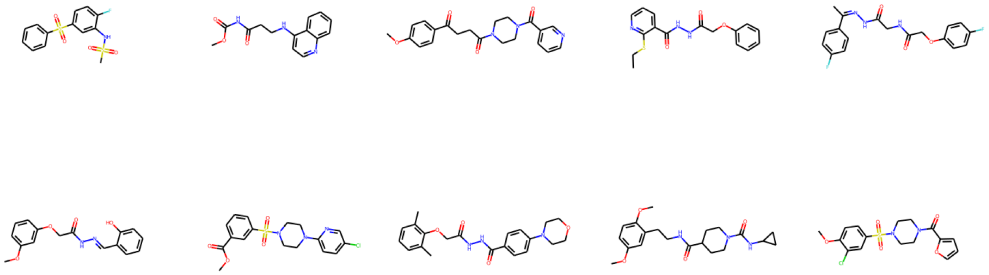

Figure 18: Example predictions of the ProbTransformer for property values of TPSA: 80.0; LogP: 2.0; SAS: 2.0 with an allowed deviation of 0.5, 0.1, and 0.1, respectively.

TPSA: 80.0, LogP: 2.0, SAS: 4.0

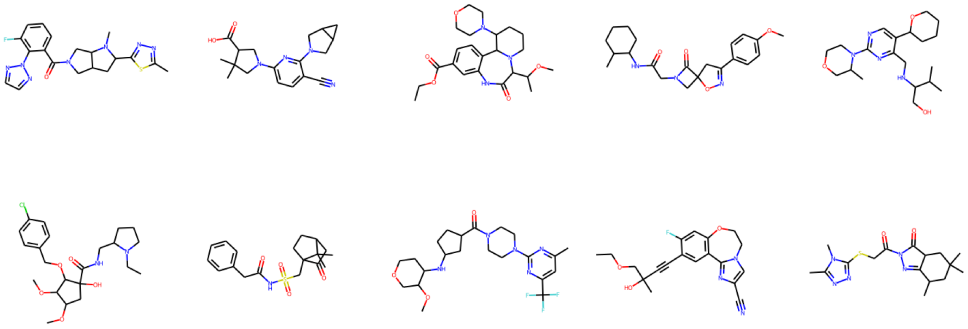

Figure 19: Example predictions of the ProbTransformer for property values of TPSA: 80.0; LogP: 2.0; SAS: 4.0 with an allowed deviation of 0.5, 0.1, and 0.1, respectively.

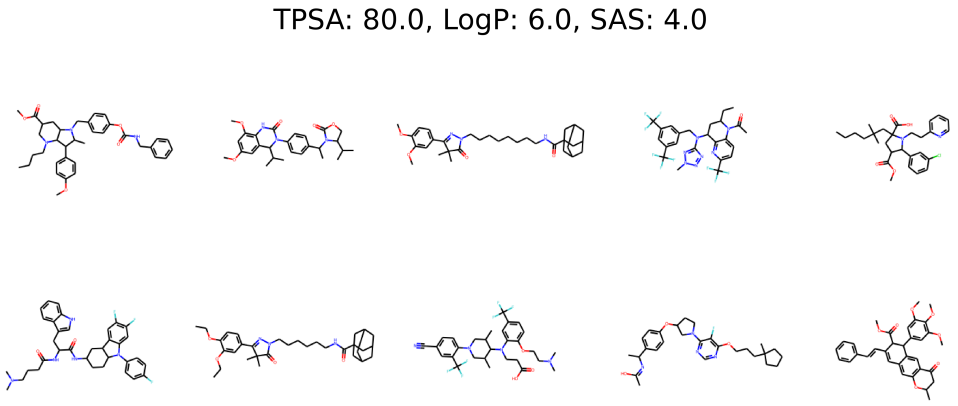

Figure 20: Example predictions of the ProbTransformer for property values of TPSA: 80.0; LogP: 6.0; SAS: 4.0 with an allowed deviation of 0.5, 0.1, and 0.1, respectively.

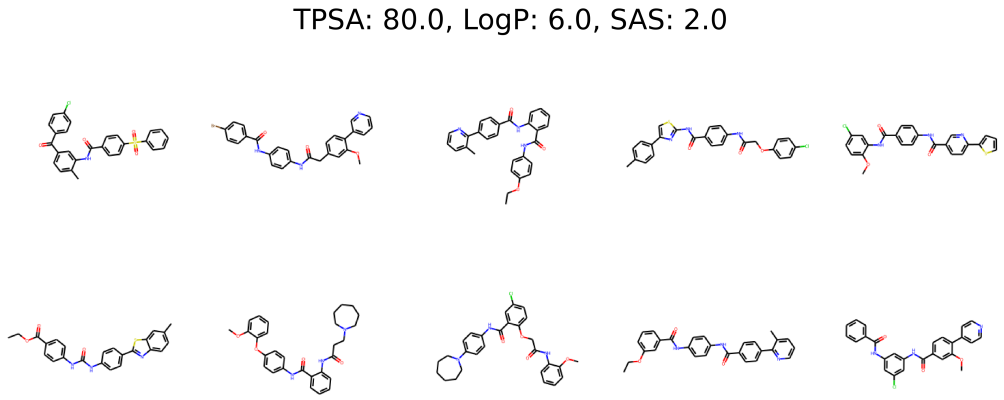

Figure 21: Example predictions of the ProbTransformer for property values of TPSA: 80.0; LogP: 6.0; SAS: 2.0 with an allowed deviation of 0.5, 0.1, and 0.1, respectively.

D.4 Related Work

Inspired by progress in the field of natural language processing (NLP), early work employed recurrent neural networks (RNNs) to produce focus libraries based on the SMILES notation [73]. Later on, these approaches were coupled with reinforcement learning (RL) to focus the generation on molecules with desirable properties [69, 74]. Additional methods were proposed to tackle the problem, including generative adversarial networks (GANs) [70, 71], variational autoencoders (VAEs) [75, 76], and adversarial autoencoders (AAEs) [72, 77, 78, 79]. For more details on the different methods, we refer the interested reader to multiple reviews of the field [83, 84, 51, 53].

More recently, the success of self-attention mechanisms entered the field and novel methods were developed, adding attention either to RNNs [16] or VAEs [16, 118]. Remarkably, Transformer-based VAEs showed more complex latent representations of molecules and outperformed previous state-of-the-art VAEs [79] in the field [16].

However, as discussed before in Section 5, only some methods yet approached the challenging task of generating molecules with (multiple) predefined property values (conditional generation) [73, 80, 81, 82, 11].

E Ablation Study

We list the hyperparameters of the Transformer and ProbTransformer as well as the training configuration for the ablation study in Table 21. We performed the training on one Nvidia RTX2080 GPU and the training time for one ProbTransformer model is ∼56h (1 prob layer) to ∼80h (all prob layer) and for one Transformer ∼33h. Also, we reduced the learning rate for the architecture ablation study due to unstable training when using all prob layers.

Table 21: Hyperparameters of the Transformer and ProbTransformer and training details of the ablation study.

Feed-forward dim	2048
Latent Z dim	512
Model dim	512
Number of layers	6
Number of heads	8
Prob layer	4,3-4,2-5,all
Kappa	0.1
Dropout	0.1
Optimizer	adamW
Beta 1	0.9
Beta 2	0.98
Gradient Clipping	100
Learning rate schedule	cosine
Learning rate high	0.0001/5
Learning rate low	0.00001/5
Warmup epochs	1
Weight decay	0.01
Epochs	200
Training steps per epoch	5000