

ASOptimizer: Optimizing antisense oligonucleotides through deep learning for IDO1 gene regulation

Gyeongjo Hwang,^{1,5} Mincheol Kwon,^{2,5} Dongjin Seo,¹ Dae Hoon Kim,² Daehwan Lee,¹ Kiwon Lee,¹ Eunyoung Kim,² Mingeun Kang,^{1,3} and Jin-Hyeob Ryu^{2,4}

¹Spidercore Inc, 17, Techno 4-ro, Yuseong-gu, Daejeon 34013, South Korea; ²BIORCHESTRA Co., Ltd., 17, Techno 4-ro, Yuseong-gu, Daejeon 34013, South Korea; ³Department of Electrical Engineering, Korea Advanced Institute of Science and Technology, 291 Daehak-ro, Yuseong-gu, Daejeon 34141, South Korea; ⁴BIORCHESTRA US., Inc., 1 Kendall Square, Building 200, Suite 2-103, Cambridge, MA 02139, USA

Recent studies have highlighted the effectiveness of using antisense oligonucleotides (ASOs) for cellular RNA regulation, including targets that are considered undruggable; however, manually designing optimal ASO sequences can be labor intensive and time consuming, which potentially limits their broader application. To address this challenge, we introduce a platform, the ASOptimizer, a deep-learning-based framework that efficiently designs ASOs at a low cost. This platform not only selects the most efficient mRNA target sites but also optimizes the chemical modifications for enhanced performance. Indoleamine 2,3-dioxygenase 1 (IDO1) promotes cancer survival by depleting tryptophan and producing kynurenine, leading to immunosuppression through the aryl-hydrocarbon receptor (Ahr) pathway within the tumor microenvironment. We used ASOptimizer to identify ASOs that target IDO1 mRNA as potential cancer therapeutics. Our methodology consists of two stages: sequence engineering and chemical engineering. During the sequence-engineering stage, we optimized and predicted ASO sequences that could target IDO1 mRNA efficiently. In the chemical-engineering stage, we further refined these ASOs to enhance their inhibitory activity while reducing their potential cytotoxicity. In conclusion, our research demonstrates the potential of ASOptimizer for identifying ASOs with improved efficacy and safety.

INTRODUCTION

Antisense oligonucleotides (ASOs) have emerged as a promising class of drugs that have the potential to target and modulate genes associated with various diseases.^{1–3} ASOs offer a unique approach to precision medicine by selectively suppressing target gene expression, which offers a promising treatment for a wide range of genetic and rare diseases.^{4–8} Despite these advantages, designing the structure of therapeutic ASOs remains a significant challenge because of the intricate interplay between RNA sequence, structure, and function.^{9–11} Conventionally, researchers have heavily relied on their molecular biology expertise along with empirical observations to generate potential ASO candidates. However, the complexity of

ASO design is evident when one considers the vast chemical space of RNA sequences, in which each position is occupied by any of the four nucleotide bases (A, U/T, G, or C). For ASOs of length l , this results in 4^l potential combinations, which becomes an astronomical number for $l \geq 20$. The complexity is further increased when chemical modifications, such as linkage and sugar modifications, are considered.

To navigate the vast search space of RNAs, researchers often resort to using ASO sequences that are fully complementary to the target gene, employing a gapmer design approach. Gapmers refer to ASOs that satisfy specific design criteria: (1) the central segment of the ASO exhibits complete complementarity to the target DNA/RNA in the antisense direction, and (2) chemically modified RNA regions, known as wings, are positioned at both ends of the sequence to enhance nuclease resistance and binding affinity.¹⁰ This simplified strategy enables one to focus on a subset of candidates that exhibit improved pharmacokinetics and toxicity profiles^{12–15} compared with plain phosphorothioate (PS) ASOs with no modifications. However, this reduced search space results in ASOs with significantly suboptimal performance in terms of both inhibition rate and cytotoxicity,^{12,13,15,16} motivating the need for a more systematic way to ASO design beyond standard gapmers. The inherent complexity of RNA mechanisms coupled with the extensive search space of RNAs requires a paradigm shift toward computational approaches.

There is a substantial body of prior work regarding the statistical analysis of therapeutic RNAs and their performance, which has primarily focused on small interfering RNAs (siRNAs). For example,

Received 17 October 2023; accepted 3 April 2024;
<https://doi.org/10.1016/j.omtn.2024.102186>.

⁵These authors contributed equally

Correspondence: Mingeun Kang, Spidercore Inc, 17, Techno 4-ro, Yuseong-gu, Daejeon 34013, South Korea.

E-mail: minkang@spidercore.io

Correspondence: Jin-Hyeob Ryu, BIORCHESTRA Co., Ltd., 17, Techno 4-ro, Yuseong-gu, Daejeon 34013, South Korea.

E-mail: branden.ryu@biorchestra.com



several groups have proposed initial protocols for the design of siRNAs,^{17–20} which were later expanded to consider off-targets and other variables.^{21,22} Based on these rules, screening algorithms were developed to filter out less-promising candidates^{23,24} and several computational models were proposed to predict the regulatory efficacy of input siRNAs.^{25–28} Moreover, in the context of exon-skipping ASOs, similar studies have been made to identify factors contributing to the regulatory performance,^{29,30} leading to the development of computational algorithms that predict exon-skipping efficacy for input ASOs.^{31,32} However, even with the extensive literature and availability of online software,^{33–37} the challenge of optimizing therapeutic RNAs with respect to chemical modifications remains unsolved. Indeed, this aspect poses one of the most crucial hurdles in designing effective therapeutic ASOs. Hence, our study addresses this problem by focusing primarily on RNase H-mediated ASOs. We propose a data-driven strategy, aiming to unlock the full potential of ASOs for precision medicine.

Our contribution lies in the development of a novel and holistic framework, the ASOptimizer, which enables the optimization of therapeutic ASOs at both the sequence and molecular level. To achieve this, we establish a two-stage approach: (1) we create a linear factor model based on early studies of microRNA (miRNA)-binding sites.^{38–42} The model characterizes the relationship between the modulation performance of ASO and input variables, such as Gibbs free energy and the secondary structure of the target, to screen out top k promising candidates. (2) We build a deep graph neural network that accepts the molecular graph of ASOs as input and predicts the regulatory efficacy, thus guiding the chemical modification of the top k ASOs to enhance their potency. To validate our approach, we conducted *in vitro* experiments using indoleamine 2,3-dioxygenase 1 (IDO1) mRNA. The protein product of this gene has an important role in cancer cell survival,^{43,44} as it has the ability to reduce tryptophan levels and produce kynurenine to suppress the immune system.^{45–48} Our results demonstrate that the proposed ASOs exhibit superior inhibition of IDO1 expression levels and lower cytotoxicity by ASO structure compared with standard gapmers.

RESULT

In this section, we show the *in silico* validation on ASOptimizer, followed by *in vitro* experiments on the recommended ASOs, with the goal of demonstrating the therapeutic potential of the proposed framework. Both *in silico* and *in vitro* results demonstrate the outstanding performance of our model compared with the currently available baselines.

Dataset for computational models

We built a database of experimental observations from granted patents and scientific papers, which included both *in vitro* and *in vivo* experiments. Patents were gathered through Lens,⁴⁹ a search platform specialized for such a purpose. Using the keyword “ASO,” we compiled more than 1,000 results and extracted tabular data from the documents. The *in vitro* data included various cell culture conditions, such as the target gene, cell line, density (cells/well), transfection,

dose amount (nM), treatment period (hours), and other relevant factors considered in the literature. In addition to the cell culture details, each sample also contained information on the ASO sequence, type/position of chemical modifications, and the inhibition rate against the target mRNA measured by real-time qRT-PCR. The database consisted of 187,090 unique samples for 67 distinct target mRNAs.

Since we aimed to validate our approach through IDO1 gene regulation, we also extracted a few samples⁵⁰ from our database in which the target was IDO1 mRNA (GenBank: NM_002164). This subset contained 155 samples. Note that this subset was used for the first stage of our approach in finding the best location for a therapeutic ASO to bind to, whereas the entire database was used for optimizing chemical modifications. For clarity in the following sections, we refer to the entire dataset as D and the subset as D_{IDO1} .

ASO-mediated gene regulation can partially be explained by thermodynamic stability and structural accessibility.

We conducted the proposed linear model and obtained the parameters that minimize the sum of squared residuals (see Equation 2 in section “materials and methods”) using the subset D_{IDO1} . As presented in Table S1, the coefficients exhibited strong statistical significance with p values of $1.20\text{e}-08$, $2.88\text{e}-02$, and $2.96\text{e}-10$ for change in Gibbs free energy after a complementary ASO binds to the target gene, average change in Gibbs free energy after the ASO binds to off-target genes, and geometric accessibility of the binding region considering the secondary structure of the target gene, respectively. This indicates the validity of our model design; especially the thermodynamic and structural features as inputs variables.

In Figures 1A and 1B, the scatterplots on experimental inhibition rates (y) versus predicted inhibition rates (\hat{y}) are presented and annotated with Pearson’s correlation. From Figure 1A, a moderate level of statistical correlation ($\rho = 0.57, \text{MSE} = 0.0435$) is evident, whereas a relatively stronger association ($\rho = 0.72, \text{MSE} = 0.0644$) is evident in Figure 1B, when considering longer sequences only. A Pearson correlation of $\rho = 0.72$ is associated with a case in which 52% of the variation in the experimental inhibition can be explained by the predicted inhibition. We found that our linear model explained the experimental data better when the input ASO was relatively longer. In Figures 1C and 1D, we also show contour plots of Pearson correlation. By observing the large slope along the y axis, the importance of considering secondary structure in the model can be seen. In contrast, the effect of considering the thermodynamic stability between the ASO and off-target genes is relatively marginal, as shown by the smaller slope along the x axis. This marginal gain becomes even less under longer sequences. We believe that our results are consistent with a recent finding that long ASOs tend to suffer less from off-target effects.⁵¹

In silico comparison with a baseline

We compared the proposed linear model with a currently available baseline: *Sfold*.³³ *Sfold* is one of the early works used for the design

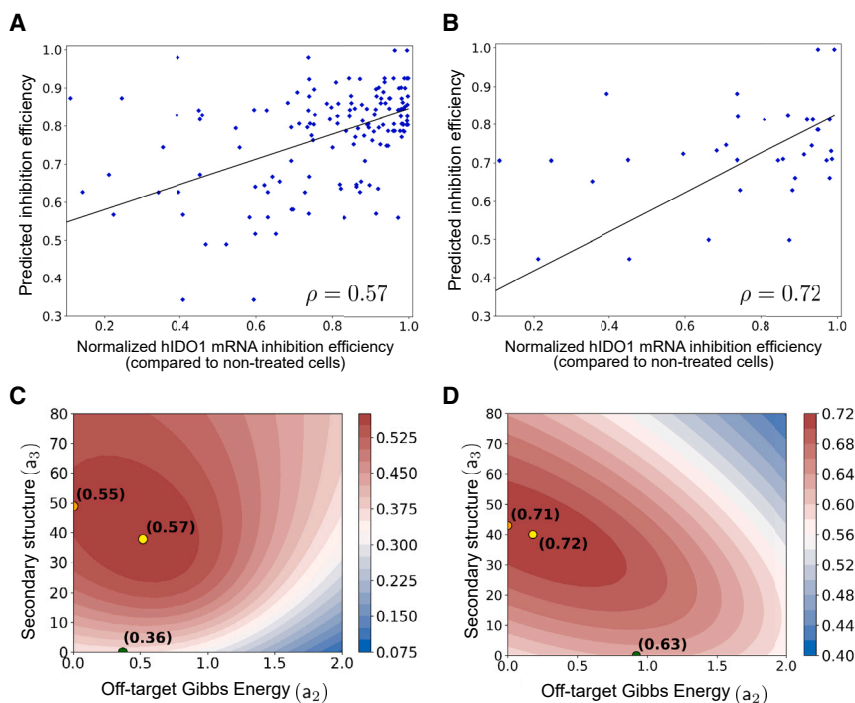


Figure 1. Graphical illustration of the result

(A and B) Experimentally observed inhibition rates (x axis) and their predicted values (y axis). ρ on the bottom right of each graph represents the Pearson correlation coefficient. In (A), we show the scatterplot using 155 samples from dataset $\mathcal{D}_{\text{IDO1}}$, while (B) employs 40 long sequences that are 17-mer or longer. In (C) and (D), we show the surface of Pearson correlation (ρ), represented by contour plots. On the x and y axes, we vary the values of normalized (divided by a_1) a_3 and a_2 (See Equation 1 in section “materials and methods”), respectively, and the correlation between model predictions and experimental inhibition rates are depicted as different colors at each coordinate. That is, we display the change in Pearson correlation due to the shift in the coefficients of our model. The yellow circles indicate the points that achieve the maximum correlation (i.e., \mathbf{a}^*). The green and orange circles indicate the x and y intercepts, each representing a case where either one of a_3 or a_2 is neglected; (C) employs $\mathcal{D}_{\text{IDO1}}$, while (D) employs only long sequences.

of siRNAs, ASOs, and trans-cleaving ribozymes, which is still available. The algorithm computes the antisense oligo binding energy using the weighted sum of DNA/RNA stacking energy for hybridization.⁵² We chose this as a single baseline to make a comparison as other baselines did not fit our setting. For instance, computational models other than *Sfold* are mainly focused on siRNAs,^{25–28} which require thermodynamic stability of the double strand itself as an input feature. On the other hand, ASO-specific models accept features such as the distance between the binding region and the closest skipping exon.^{31,32} Note that such features are not considered in our problem context, where the primary focus is RNase H-mediated ASOs. Hence, computational models for neither the siRNAs nor exon-skipping ASOs are appropriate as baselines.

We split $\mathcal{D}_{\text{IDO1}}$ into training and test sets at an 8:2 ratio and found linear model parameters (see Equation 2 in section “materials and methods”) using the training set. We then compared our model with the baseline on the remaining test set. We found that our model achieved Pearson correlation and mean squared error (MSE) of $\rho = 0.66$ and $\text{MSE} = 0.0297$, respectively, while *Sfold* achieved $\rho = 0.50$ and $\text{MSE} = 0.0361$. We believe that the results are plausible, given that our model considers more factors compared with *Sfold*, such as off-targets and secondary structure.

In silico validation through external data

To show the generalizability of our linear model, we also conducted a simple *in silico* study as follows: we prepared a dataset consisting of experimental results on PLP1, APOL1, and IDO1. We used PLP1 and APOL1 samples as training data while IDO1 samples were used

as test data. In this case, no information about IDO1 can be inferred from training samples. Under this external validation setting, we found that our model attained $0.6016 (\pm 0.1326)$ Pearson correlation and $0.0356 (\pm 0.0086)$ MSE on the IDO1 test set. Compared to test performance using internal samples as training data, the compromise in performance was relatively modest, proving the robustness of our linear model.

The trained model produces rankings of complementary ASO sequences based on the predicted IDO1 inhibition

Using the trained linear model, we computed the scores for candidate ASOs of 19 bases in length that were complementary to the IDO1 mRNA. Note that the sequence length is a design parameter that is determined by bio-scientists, as in many other software tools available.^{33–37} Our choice of 19 bases in length was based on our familiarity with ASOs of length 19–22, and our finding that model predictions are more consistent with experimental data on longer sequences. Also, we set the number of considered off-target genes to 50, in contrast to the 10 used in the *in silico* analysis, in order to account for more off-target genes in the *in vitro* validation (see x_2 in Equation 1 in section “materials and methods”). We argue that our model is not sensitive to the number of off-target genes parameter as shown by a simple ablation study in Figure S1. One can use any values less than 500, since the prediction performance does not change significantly.

Using this setting, we identified six ASO candidates that exhibited the highest prediction scores in the population of complementary sequences of 19 bases in length. As presented in Figure S2, the top six candidates are positioned at the right end of the distribution, over $\mu + 2\sigma$. These examples were selected for further testing in *in vitro* experiments. More detailed information is listed in Table S2.

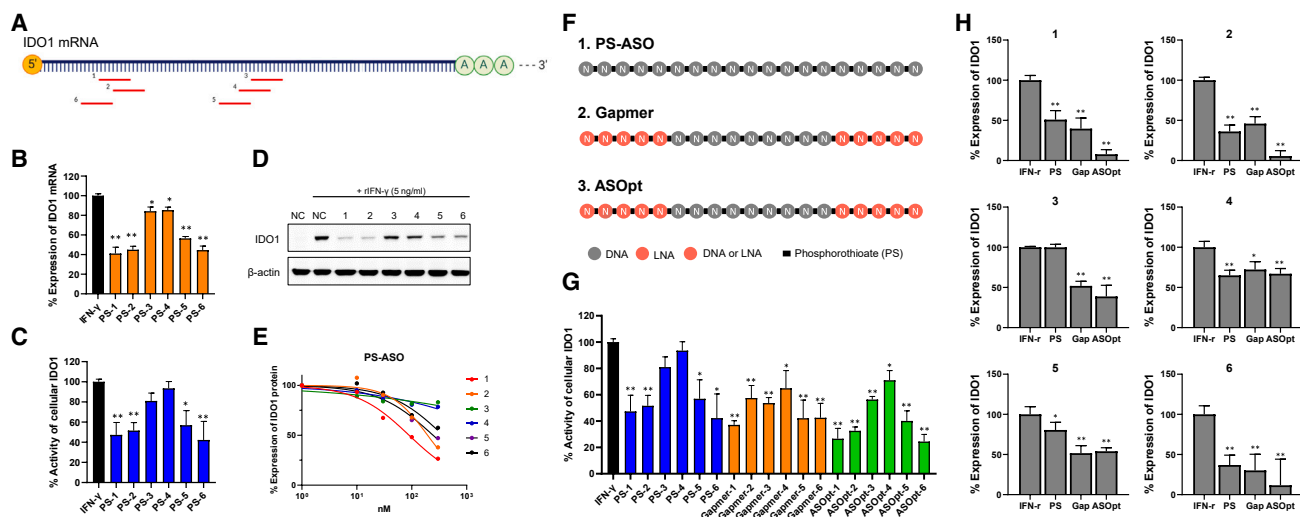


Figure 2. Performance comparison of ASOpt sequences versus PS-ASOs and gapmers

(A) Visualization of targeting regions of the six PS-ASOs. There are notably two distinct locations in the mRNA; one near the 5' end and the other in the coding sequence (CDS). (B) The qRT-PCR results indicate that the expression of IDO1 mRNA was reduced by 300 nM PS-ASOs after 48 h in HeLa cells. (C) Kynurenine production resulting from reduced cellular IDO1 activity by 300 nM PS-ASOs after 48 h in HeLa cells. (D) WB result indicating that the expression of IDO1 was reduced by 300 nM PS-ASOs after 48 h in HeLa cells. (E) A dose-dependent reduction in the expression of IDO1 was determined by WB. (F) Schematic diagram of the structures of PS-ASO, gapmer, and ASOpt. (G) Kynurenine production resulting from reduced cellular IDO1 activity by 300 nM PS-ASO (PS), gapmer (Gap), and ASOpt after 48 h in HeLa cells. (H) Normalized graphs versus β - actin. WB results indicating that the expression of IDO1 was reduced by 300 nM of PS-ASO, gapmer, and ASOpt for 48 h in HeLa cells (We conducted three independent experiments and drew the bar graphs with error bars using means and standard deviations, * $p < 0.05$, ** $p < 0.01$ vs. rIFN- γ -treated negative control).

The six ASO sequences recommended by ASOptimizer were synthesized using the PS backbone. Named after the backbone, we refer to them as PS-ASO-1 through PS-ASO-6. We first determined whether they could reduce IDO1 mRNA and the corresponding protein at 300 nM in HeLa cells. The targeting sites of PS-ASOs in IDO1 mRNA are illustrated in Figure 2A. Figures 2B and 2C showed the inhibitory effects on IDO1 from the qRT-PCR and cellular kynurenine production. All PS-ASOs successfully reduced IDO1 expression compared with the recombinant interferon (rIFN)- γ -treated negative control. In particular, PS-ASOs 1, 2, 5, and 6 exhibited superior modulation performance to PS-ASOs 3 and 4, with higher statistical significance. To confirm the reduction of IDO1 protein expression we performed western blot (WB) assay. As shown in Figure 2D, PS-ASOs 1, 2, 5, and 6 resulted in a large decrease in cellular IDO1 expression. In Figure 2E, the inhibition performance at different concentrations was also identified (10, 30, 100, 300 nM). We display the dose-response curve, in which intermediate data points were interpolated using GraphPad Prism 10. Half-maximal inhibitory concentration (IC_{50}) values for each PS-ASO were also estimated using the same software listed in Table S3. These results demonstrate that the ASOs derived from ASOptimizer effectively target IDO1.

In silico validation on the chemical-engineering module suggests that ASOptimizer can improve upon PS-ASOs

As previously described, we generated training and test pairs, $D_{\text{pairs}}^{\text{train}}$ and $D_{\text{pairs}}^{\text{test}}$, respectively, for the Edge-augmented Graph Transformer (EGT)⁵³-based deep neural network in the chemical-engineering

phase. For *in silico* validation, we measured the predictive accuracy of the model trained on $D_{\text{pairs}}^{\text{train}}$ against the remaining test set $D_{\text{pairs}}^{\text{test}}$ (accuracy in predicting the superior one from a pair of positive and negative samples). This resulted in a 73.39% accuracy, indicating the considerable potential for identifying which chemical modification is better on unknown examples. We further compared our chemical-engineering module with a currently available baseline.¹³ See the Figures S3 and S4 for an in-depth analysis of the results.

Next, we applied the trained model to find more preferred chemical modifications to the six PS-ASOs found in the sequence-engineering phase. Here, we constrained the space of chemical variants to PS and locked nucleic acids (LNAs), since the majority of the samples in our database use only a single type of chemical modification in a sequence and LNA accounts for the largest portion in chemical modifications. The reduced chemical space also helps us simplify the *in vitro* validation process, decreasing the unnecessary complexity from the experiments. Similar to the previous stage, we tested the model with all possible PS-vs-LNA combinations, resulting in $2^4 = 16$ cases for each wing of five nucleotides (RNAs at positions 2–5 were modified, respecting the design rule in related work¹³). This amounts to a total of $2^8 = 256$ distinct LNA modification patterns for each sequence. From the ranking of the predicted scores, we identified the modification with the highest score for each of the PS-ASOs (see Figure S5). In Table 1, we summarize the sequence information of the optimized ASOs. For a graphical abstraction of the sequence information, see Figure 2F.

Table 1. The top six ASO sequences obtained through the sequence-engineering module

Index	PS-ASO	Gapmer	ASOpt
1	caaggcgtgtgacttg	CAAGGcgetgtgacTTGTG	CAAGcgtgtgacttGTG
2	ctagacgtgcaaggcctg	CTAGAcgtgcaaggCGCTG	CTAgacgtgcaaggcGCTG
3	caaacactcaggactgaggg	CAAACtcaggactGAGGG	CaAActcaggactgaGGG
4	cggactgagggattgact	CGGACtgaggattTGACT	CgGActgagggattgACT
5	actcaggactgagggatt	ACTCACggactgagGGATT	ACTCacggactgagggATT
6	cgctgtgacttggctg	CGCTGtgacttggGTCTG	CGCTgtgacttgggtCTG

Bold and uppercase, LNA; lowercase, DNA.

PS-ASO denotes phosphorothioate antisense oligonucleotides, gapmer has LNA modifications on both 5' and 3' ends (5-9-5), and ASOpt represents chemically optimized ASOs derived from the chemical-engineering module of ASOptimizer.

***In vitro* results prove that ASOptimizer indeed improves the PS-ASOs via optimized chemical modifications**

We synthesized the six chemically optimized ASOs (named ASOpt) found above and tested their performance *in vitro*. As shown in Figure 2G, we compared the cellular IDO1 enzyme activity measured by kynurenine levels in HeLa cells between the six PS-ASOs, 5-9-5 gapmers, and ASOpt-1 through ASOpt-6. Here, PS-ASO-*i*, gapmer-*i*, and ASOpt-*i* share a common nucleotide sequence but have distinct chemical modifications. As shown in the bar graph in Figure 2G, gapmers tended to outperform their PS-ASO counterparts, with gapmer-2 and -6 representing exceptions to the trend. Interestingly, the ASOpt sequences exhibited even stronger inhibitory effects compared with the gapmers. This trend was more evident for ASOpt-1, ASOpt-2, and ASOpt-6, whereas the others did not show a strong improvement over the gapmers. A similar observation was made in experiments to test the expression levels. WB and qRT-PCR results indicated that the expression of IDO1 was decreased significantly for ASOpt-1, ASOpt-2, and ASOpt-6 (Figures 2H and S6A). We drew the same conclusion from experiments conducted on U87-MG cells. See Figure S6B for the results on U87-MG cells, in which we also observed even better improvements over the PS-ASOs and gapmers.

One noteworthy aspect of the experiment was that the PS-ASO-3 appeared ineffective, suggesting that the binding site found by our sequence-engineering module may not be working. However, after applying chemical modifications, both gapmer-3 and ASOpt-3 suppressed the IDO1 expression. This highlights the importance of optimizing a therapeutic RNA at the molecular level.

We also generated the dose-response curves for the candidates, as shown in Figure 3A. Similar to above, ASOpt-1, 2, and 6 exhibited the most noticeable improvement compared with the PS-ASOs and gapmers with respect to the dose-dependent inhibitory performance. ASOpt-3, which initially showed no activity in the PS-ASO, exhibited remarkable inhibitory activity upon chemical modification. As shown in Table 2, the IC₅₀ values determined by GraphPad Prism 10 and the corresponding fold changes are presented. We observed an approximate 6.19-fold improvement for ASOpt-3 against PS-ASO-3, and a 4.02-fold change for ASOpt-3 against gapmer-3. The other candidates showed different scales of fold values, yet mostly over the threshold of

1 (indicating improvement), except for sequence 5. These findings confirmed that our holistic framework ASOptimizer has the potential to find effective ASO sequences and their chemical modifications.

Enhanced cell viability with chemically optimized ASOs

Another intriguing yet unexpected observation was that the cytotoxicity levels were markedly decreased with all chemically optimized ASOs compared with PS-ASOs and gapmers. Specifically, we administered 1,000 nM PS-ASOs, gapmers, and ASOpt to HeLa cells for 48 h, and measured lactate dehydrogenase (LDH) levels. As shown in Figure 3B, PS-ASOs exhibited the strongest cytotoxicity in all six candidates and was in contrast to the ASOpt-1 through ASOpt-6, which showed reduced cytotoxicity of more than 50%. Because our chemical-engineering module was not trained to estimate the cytotoxicity of an input ASO, but inhibitory efficacy, we could not expect such improvement. One probable interpretation of the result is that, as our model prediction removes one or more LNA-modified nucleotide bases in each wing, toxicity is enhanced due to reasons explained in a recent study.¹⁵ An in-depth analysis on the results is presented in the “discussion” section.

Chemically optimized ASOs improve THP-1 M1 differentiation

Macrophages are immune cells capable of differentiating into either M1 or M2 phenotypes based on the surrounding environment. M1 differentiation is induced by internal cytokines, such as rIFN- γ and tumor necrosis factor (TNF)- α , as well as external pathogens, such as lipopolysaccharide (LPS).^{54,55} However, the presence of IDO1 inhibits M1 differentiation, even in response to rIFN- γ and LPS stimulation.⁵⁶ Since M1 macrophages function as killer cells when targeting cancer, the induction of M1 differentiation is one of the potential strategies for developing anti-cancer drugs.⁵⁷ This prompted us to determine whether the ASOs discovered by ASOptimizer can reduce IDO1 expression in THP-1 macrophages, resulting in improved M1 differentiation. A schematic representation of the experimental procedures is depicted in Figure 4A. In M1-differentiated THP-1 cells, IDO1 mRNA was significantly suppressed by gapmer and ASOpt sequences compared with PS-ASO treatment (Figure 4B). Conversely, the expression of M1 macrophage markers, TNF- α and CXCL10 mRNA, exhibited a proportional increase in response to the reduction of IDO1 expression (Figure 4C). These findings indicate that ASOs

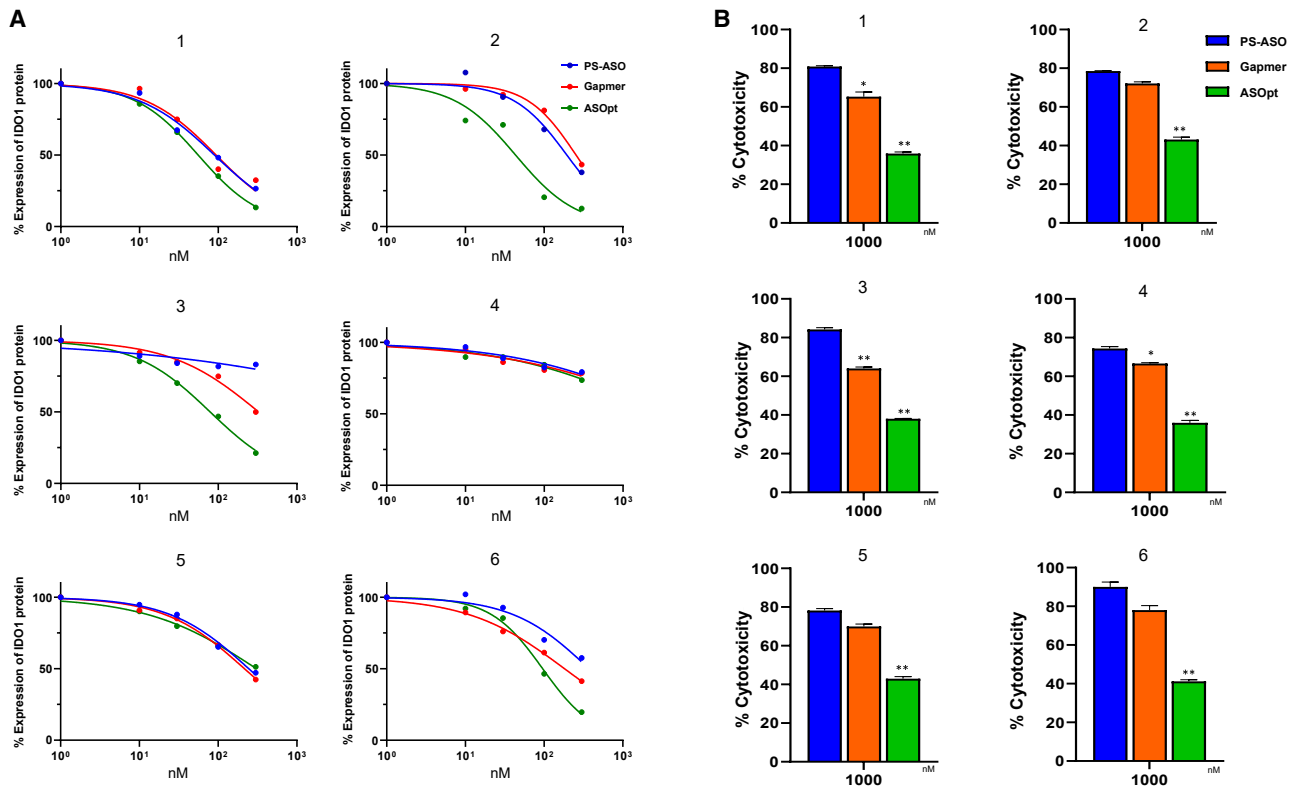


Figure 3. Dose-dependent and cytotoxicity analysis

(A) Dose-dependent analysis for comparative evaluation and IC_{50} determination of ASO sequences. These results were obtained by treating HeLa cells with PS-ASO, gapmer, and ASOpt at concentrations of 10, 30, 100, and 300 nM for 48 h, followed by quantification of IDO1 protein levels. Red line, PS-ASO; yellow line, gapmer; green line, ASOpt. (B) Reduced cytotoxicity following ASOpt treatment compared with PS-ASO and gapmer in HeLa cells. The results were obtained by treating HeLa cells with PS-ASO, gapmer, and ASOpt at concentration of 1,000 nM for 48 h (Bar graphs with error bars were drawn using means and standard deviations from three independent experiments, * $p < 0.05$, ** $p < 0.01$ vs. PS-ASO).

discovered by ASOptimizer can effectively reduce IDO1 expression in immune cells, thereby promoting M1 differentiation. This supports their potential use in inducing cancer-killing M1 macrophages.

DISCUSSION

We proposed a holistic framework for optimizing RNase H-mediated ASOs, not only at the sequence level but also at the molecular level. The proposed platform, ASOptimizer, consists of two sequential components: (1) a sequence-engineering module based on a linear factor model that relates the regulatory performance of an ASO to its thermodynamic features and secondary structure, and (2) a chemical-engineering module based on a deep graph neural network that predicts the regulatory performance of an input ASO. Primarily focusing on suppression of the IDO1 mRNA, the biological relevance of ASOptimizer was validated through multiple experiments. Here, we briefly summarize the two components along with detailed discussions and future research directions.

Sequence engineering

We trained the proposed linear model using a dataset D_{IDO1} of experimental inhibition rates against the IDO1 gene, which is a sub-

set of the entire database D . The trained linear model was used to sort complementary ASOs of 19 bases in length, resulting in six candidates positioned at the right end of the histogram of prediction scores (see Figure S2). ASOs 1, 2, and 6 target locations near the 5' end, while the others target the CDS region (Figure 2A). These sequences with a PS backbone, PS-ASOs, were tested *in vitro* and their effectiveness in modulating the IDO1 expression was confirmed (Figures 2B–2E).

The linear model can be improved

Although our model parameters achieved statistically significant p values (Table S1), the overall goodness-of-fit was not satisfactory: a Pearson correlation of 0.72 for longer sequences only. More complex models that are possibly non-linear or more input features are needed to better describe the input ASO as in previous studies of siRNAs and exon-skipping ASOs. We also considered such extensions as presented in the supplemental information (Tables S4–S6), although it was not easy to confirm a significant performance improvement. However, we believe that considering a wider range of features with increased dataset size will bring advantage for more complex models.

Table 2. IC₅₀ and fold changes between the ASOs targeting IDO1 mRNA

Index	IC ₅₀			Fold change	
	PS-ASO (nM)	Gapmer (nM)	ASOpt (nM)	PS-ASO/ ASOpt	Gapmer/ ASOpt
1	84.9	71.6	69.9*	1.21	1.02
2	213.1	277.9	53.2*	4.01	5.22
3	556.3	360.9	89.8*	6.19	4.02
4	1,450.9	1,782.4	943.6*	1.54	1.89
5	249.0	241.4*	396.9	0.63	0.61
6	346.8	221.3	94.1*	3.69	2.35

The asterisks indicate the best results among PS-ASO, gapmer, and ASOpt.

Another direction to enhance the linear model is to use more advanced and sophisticated estimation models for the input features. For example, instead of using miRanda,³⁸ one may resort to other algorithms that produce more accurate predictions of Gibbs free energy. This applies to the prediction of secondary structure as well, in which the Zuker-style algorithm⁵⁸ may be replaced by modern deep-learning-based predictive models.^{59–61} In particular, for RNase H-mediated ASOs, we stress the need to utilize more robust and accurate structure prediction algorithms to enhance the reliability of the current model. Indeed, we display the change in Pearson correlation resulting from the different secondary structure prediction algorithms in [Figure S7](#).

Finally, a feature that we can introduce into the linear model is the possibility of incorporating non-complementary ASOs. We narrowed the search space of nucleotide sequences down to a set of complementary sequences against IDO1 mRNA; however, there may be many ASO sequences that are not fully complementary yet have greater therapeutic potential in the entire space of 4^l combinations. One can also use a different optimization scheme that utilizes genetic algorithms or reinforcement-learning algorithms, which have been shown to perform well in discrete spaces.

Chemical engineering

We trained an EGT⁵³-based graph neural network using the pairwise loss function (see [Equation 5](#) in section “[materials and methods](#)”), to avoid the need of incorporating environmental variables, such as transfection, dose, and cell line. Using the trained model, we sorted the predicted results of all possible LNA modifications on the six ASOs to screen out the top performer, resulting in the ASOpt-1 through ASOpt-6 ([Table 1](#)). Experiments demonstrated that chemically optimized ASOs can outperform PS-ASOs and gapmers, which are the conventionally used baselines ([Figures 2](#) and [3A](#)). The performance increase was most noticeable for ASOpt-1, 2, and 6. Furthermore, we also observed that all six ASOpt sequences induced significantly less cytotoxicity ([Figure 3B](#)). Although the precise reason remains to be determined, we suspect that it is due to a reduction in LNA modifications, as a reduction in LNA yields decreased cytotoxicity.¹⁵ ASOpt sequences were also found

to improve THP-1 M1 differentiation, indicating their potential to activate cancer-killing M1 macrophages ([Figure 4](#)).

Future extensions to the chemical-engineering module

We can extend our model to predict other pharmacokinetic properties. For example, once we modify the loss function to include an absorption ratio prediction term, we can obtain a model that provides rankings based on modulation *and* absorption performances. This is relatively a straightforward extension of our model, but it will markedly increase the usefulness of our second module.

The fact that our experiments were limited to the LNA modifications also suggests a future research direction. Since our model was trained with a database containing various chemical modifications, the model has the capability, at least computationally, of providing meaningful recommendations on the variety of modifications; however, such a capability must be thoroughly tested by biological experiments in the future. Besides the sugar modifications, one can also explore additional possibilities, such as linkage modifications (e.g., PS, phosphodiester, and stereochemistry), to broaden the scope of our approach further. These extensions must be accompanied with well-designed *in vitro* experiments and even *in vivo* experiments where possible.

In conclusion, our novel framework, ASOptimizer, successfully provides ASO sequences with non-trivial chemical modifications, highlighting the potential of this computational approach to optimal ASO design.

MATERIALS AND METHODS

The proposed platform, ASOptimizer, consists of two sequential components: (1) sequence engineering and (2) chemical engineering. For the former, we used a linear factor model that captures the relationship between several features of an ASO and its modulation efficacy. This model was able to guide us in choosing the best possible binding sites on the target mRNA, thus resulting in candidate ASOs with complementary sequences. Then, using a deep graph neural network, we modified the candidate ASOs to improve their regulatory performances. The parameters of the two models were learned from experimental data gathered from patents and scientific publications. An overview of the entire process is presented in [Figure 5](#). In this section, we explain the details of data gathering, preprocessing, the two computational models, and our experimental methodology.

Database construction

Building an appropriate database is always the first step toward a successful learning-based model. Hence, we mainly gathered published data samples of *in vitro* experiments from [Lens.org](#). Specifically, we conducted the following iterative steps:

- (1) Selection of genes. Initially, we screened which genes to collect in our database. This step was essential because data gathering is highly labor intensive. Our main criteria here was how many relevant documents with a sufficient number of experimental samples could be found.

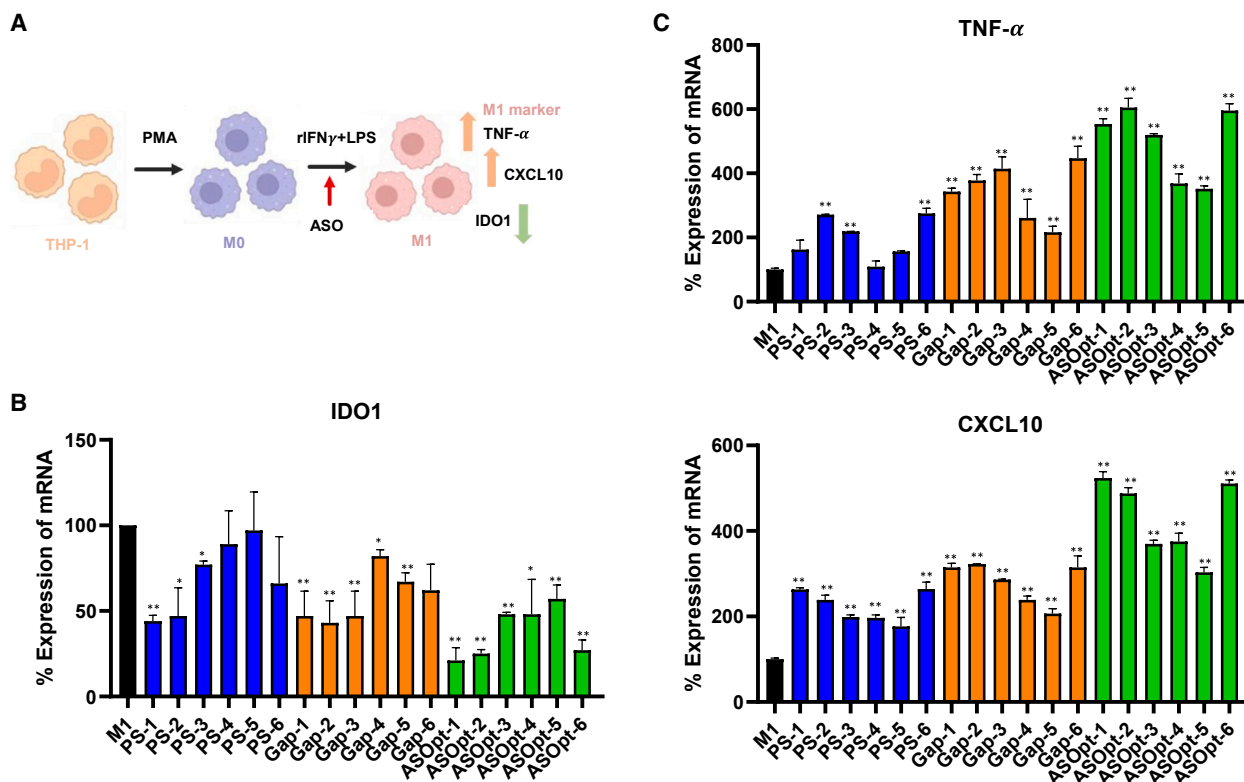


Figure 4. Increased M1 differentiation by ASOpt treatment compared with PS-ASO and gapmer in THP-1 cells

(A) Changes in THP-1 M1 differentiation and gene expression patterns after ASO treatment. (B) qRT-PCR results indicating that the expression of IDO1 mRNA was reduced by PS-ASO, gapmer, and ASOpt at concentrations of 300 nM after 48 h in M1-differentiated THP-1 cells. (C) qRT-PCR results showing that the expressions of TNF- α and CXCL10 mRNA were increased by PS-ASO, gapmer, and ASOpt at concentrations of 300 nM for 48 h in M1-differentiated THP-1 cells (Bar graphs with error bars were drawn using means and standard deviations from three independent experiments, * $p < 0.05$, ** $p < 0.01$ vs. M1-differentiated negative control).

- (2) Search in Lens.org. To access a diverse range of relevant patent documents and publications, we mainly utilized [Lens.org](https://www.lens.org/). Using this platform, we conducted targeted searches with specific keywords related to gene modulation, such as “modulators,” “modulation,” “expression,” or “compounds,” in conjunction with the selected target genes. For instance, we might use a query “modulators of IDO1.”
- (3) Data extraction from patent documents and publications. Once we identified pertinent documents, we systematically extracted relevant information; especially tabular data containing nucleotide sequences and inhibition performance metrics. The extraction process could be carried out either manually or programmatically, depending on the availability and format of the data.
- (4) Extraction of additional experimental details. In addition to sequence and inhibition data, we also extracted [supplemental information](#), such as details about chemical modifications applied to ASOs and experimental conditions. This encompassed information about cell lines, culture conditions, and other relevant experimental parameters.
- (5) Repetition. We repeated steps (1)–(4) until we had a sufficient number of data samples to train a deep neural network. Collecting more data samples was crucial to ensure that our database was

comprehensive and representative of the various ASOs and their modifications.

To obtain the final database, five researchers spent a few hundred man-hours throughout the 3 months. This laborious and demanding task served as a stepping stone toward a learning-based optimization of ASOs and their chemical modifications.

Sequence engineering

To identify the optimal binding sites for ASOs, we propose a linear factor model that relates the gene expression level with three input parameters: (1) change in Gibbs free energy after a complementary ASO binds to the target gene, (2) average change in Gibbs free energy after the ASO binds to off-target genes, and (3) geometric accessibility of the binding region considering the secondary structure of the target gene. The reasoning behind the selection of these features is based on previous studies on miRNA-binding sites. The importance of sequence complementarity and thermodynamic stability of the RNA-target heteroduplex has been repeatedly emphasized in the literature.^{38–42} Therefore, we took ΔG into account as the first feature of our approach. Taking one step further, we considered the thermodynamic stability of the heteroduplex hypothetically formed

1. Database

Table of Experimental Results

No	Target Gene	ASO Sequence (5'-3')	Modification	Inhibition Rate
1	A	GAGGCGGACACTGATTGA	cEt	82%
2	A	GAGGCGGACACTGATTGA	cEt,oMe	90%
3	B	CACGGTTTAGCAGAAGGT	MOE	54%
⋮		⋮		⋮
m-n+1	IDO1	CATCAAGCTTCCCCTTGAT	LNA	57%
⋮		⋮		⋮
m	IDO1	TGGCAGAGAGGACGCCTTC	LNA	78%

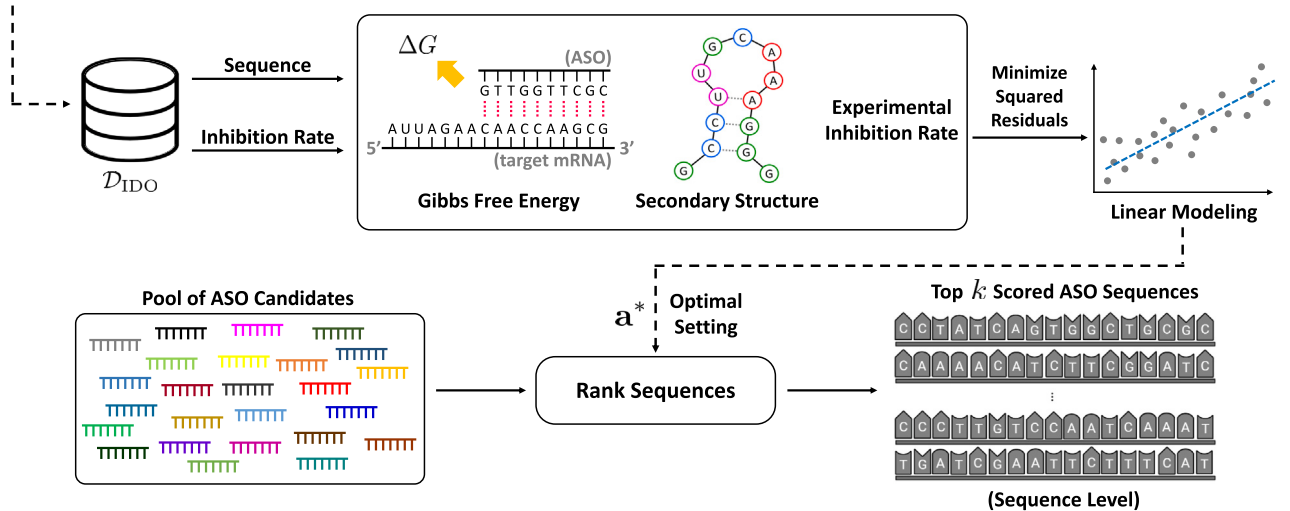
Pairwise
Matching

$$\mathcal{D}_{\text{pairs}} = \{(x_{\text{pos}}^{(i)}, x_{\text{neg}}^{(i)})\}_{i=1}^{m_{\text{pairs}}}$$

Extract
IDO Samples

$$\mathcal{D}_{\text{IDO}} = \{(x^{(i)}, y^{(i)})\}_{i=1}^n$$

2. Sequence Engineering



3. Chemical Engineering

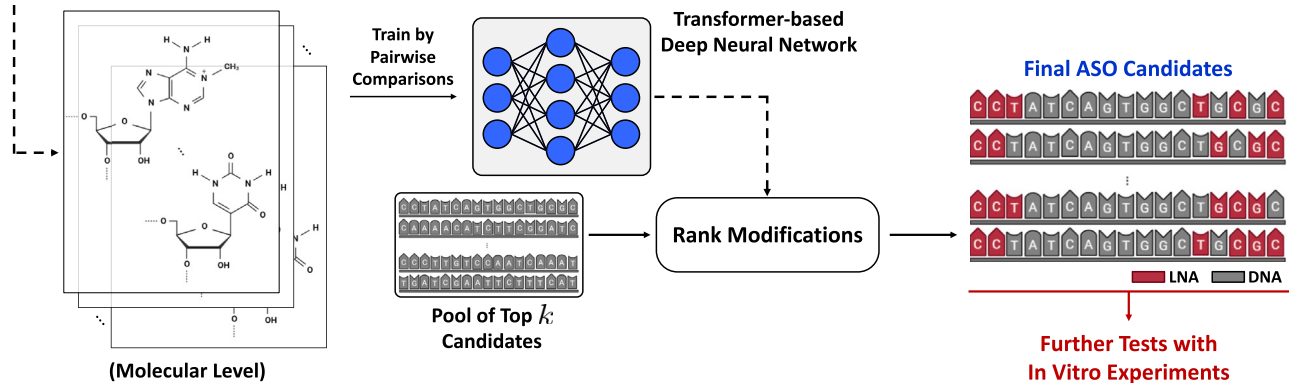



Figure 5. A schematic representation of ASOptimizer

The framework consists of a database and two computational models. (1) Database sourced from 187,090 samples of experimental observations from US patents. Within the entire dataset, denoted as \mathcal{D} , we retrieve a subset $\mathcal{D}_{\text{IDO1}}$ that contains ASOs targeting the IDO1 gene. (2) Sequence engineering that trains a linear model using the filtered dataset $\mathcal{D}_{\text{IDO1}}$ to suggest candidate ASO sequences capable of regulating the target mRNA at a sequence level. (3) Chemical engineering deepens the investigation to a molecular level using the entire dataset \mathcal{D} and pairwise matching. In the chemical-engineering phase, a transformer-based deep neural network is trained to identify chemical modifications with superior regulatory performance.

A Training Phase



B Inference Phase




Figure 6. Visualization of the chemical-engineering module's functionality

The ASO sequences are given as inputs. DNAs are indicated with black letters and LNAs are in red. The sequences are transformed into SMILES strings and then reshaped into molecular graphs, which are provided as inputs to a graph transformer model. (A) Training phase: the model $f_{\theta}(\cdot)$ is trained with a pair of ASOs (positive and negative instances) that share the same nucleotide sequence but differ in chemical modifications. The model is trained to assign higher values for the positive instance over the negative counterpart. (B) Inference phase: for a given nucleotide sequence, the trained model generates all possible combinations of chemical variants and evaluates the rankings of the combinations. The ranking information is used to identify the pattern with the highest score.

with off-target genes. Because there may be multiple off-targets, we took the average of the Gibbs free energy change over all possible cases of heteroduplexes formed with different off-target genes. Also, we took the secondary structure of the target gene into account since locally open regions are more accessible than closed and hybridized regions for therapeutic RNAs.⁶² Instead of this simple model, one may consider using more complex models. Our choice of a linear function as a regression model is mainly due to the lack of samples for the IDO1 gene. In such data-scarce scenarios, complex models run the risk of overfitting, while simpler models perform more robustly. We found that this was the case. See the [supplemental information](#) for a detailed explanation and ablation studies on this.

More formally, let $\mathcal{D}_{\text{IDO1}} = \{(x^{(i)}, y^{(i)})\}_{i=1}^n$ be our dataset of n examples, where $x^{(i)} = (1, x_1^{(i)}, x_2^{(i)}, x_3^{(i)}) \in \mathbb{R}^4$ represents the four dimensional feature vector of the i -th ASO sequence and $y^{(i)} \in \mathbb{R}$ represents the *in vitro* knockdown efficacy of the i -th ASO sequence against the IDO1 gene. Note that the constant value in the first dimension $x_0 = 1$ introduces a bias term into the model. We let x_1 be the difference in Gibbs free energy ΔG_{target} before and after the formation of the ASO-target heteroduplex computed from an off-the-shelf software, miRanda.³⁸ For the second feature x_2 , using the same software and a list of mRNAs from the RefSeq database,⁶³ we simply computed the statistical average of the Gibbs free-energy change $\mathbb{E}[\Delta G_{\text{off-target}}]$ over a set of chosen off-target genes. The set of off-target genes consisted of 10 mRNAs from the RefSeq database that showed the largest change in Gibbs free energy with the input ASO. Finally, we defined and computed x_3 as the ratio of non-hybridized nucleotides in the

binding region of the target mRNA over the length of ASO, given the predicted secondary structure of the target. We employed a Zuker-style algorithm^{58,64} to produce the structure predictions.

In this setting, the proposed linear model is:

$$\hat{y} = a_0 + a_1x_1 + a_2x_2 + a_3x_3 \quad (\text{Equation 1})$$

or equivalently, $\hat{y} = \mathbf{a}^T \mathbf{x}$ for $\mathbf{a} = (a_0, a_1, a_2, a_3)$. Here the coefficients a_0, \dots, a_3 determine the relative importance between different input features. The optimal model parameter $\mathbf{a}^* = (a_0^*, a_1^*, a_2^*, a_3^*)$ was determined by minimizing the sum of squared residuals as in standard linear regression:

$$\mathbf{a}^* = \underset{\mathbf{a}}{\operatorname{argmin}} \sum_{i=1}^n (\widehat{y}^{(i)} - y^{(i)})^2 \quad (\text{Equation 2})$$

Using our prediction model with \mathbf{a}^* , we could prioritize the complementary ASOs, resulting in the top k candidates.

Chemical engineering

Next, we further engineered the candidate ASOs to obtain better inhibition efficacy through deep learning. As shown by a few recent works, even a tiny structural difference in an oligonucleotide, such as the choice of backbone modification, results in a significant change in pharmacological properties.^{12,13} Thus, we propose a computational task of predicting modulation performance, in which a model outputs

the score of gene regulation for an input ASO. This type of model can be used to prioritize ASO candidates with different chemical modifications according to their inhibition efficacy and possibly cytotoxicity. See Figure 6 for a graphical overview of our chemical-engineering module. In the module, we considered chemical variants, including LNAs, 2'-O-ethyl, 2'-OMe, and 2'-F base modifications. To illustrate our deep-learning approach, we first describe our data-preprocessing method and our neural network architecture. Then, we discuss the training and inference procedures along with the implementation details.

Data preprocessing and neural network architecture

First, we converted our ASO data into molecular graphs. A molecular graph \mathcal{G} is formally defined as a pair of two sets: a set of atoms \mathcal{V} (V from vertices) and a set of bonds between two atoms \mathcal{E} (E from edges). Here, the set of bonds \mathcal{E} is a subset of the Cartesian product $\mathcal{V} \times \mathcal{V}$. As chemical bonds are undirected, molecular graphs are treated as undirected graphs; if $e_{ij} = (v_i, v_j) \in \mathcal{E}$, then $e_{ji} = (v_j, v_i) \in \mathcal{E}$. One equivalent representation of \mathcal{E} involves a binary symmetric matrix known as the adjacency matrix $A \in \{0, 1\}^{N \times N}$, where N is the number of atoms in the molecule. In the adjacency matrix, each entry A_{ij} holds the value 1 if $(v_i, v_j) \in \mathcal{E}$, and 0 otherwise.

To obtain a proper input representation for the graph neural network, we first transformed the ASO sequences into the Simplified Molecular-Input Line-Entry System (SMILES) format using an online software, ChemDraw.⁶⁵ Then, the Python library RDKit⁶⁶ enabled us to change the SMILES strings into molecular graphs. Finally, following the approach introduced by Open Graph Benchmark (OGB),⁶⁷ we retrieved relevant feature vectors of dimensions 9 and 3 for the atoms and bonds, respectively, together with the adjacency matrix of the graph. As a result, we obtained a dataset $\mathcal{D} = \{x^{(i)}, y^{(i)}\}_{i=1}^m$ of m examples for chemical engineering, where each $x^{(i)} = (\mathcal{V}^{(i)}, \mathcal{E}^{(i)}, A^{(i)})$ was the triplet for the i -th ASO and $y^{(i)}$ was the experimentally observed inhibition rate.

For the neural network architecture, we employed Edge-augmented Graph Transformer,⁵³ a transformer-based model,⁶⁸ which achieves state-of-the-art performances in multiple graph-related tasks as shown by recent works.^{69,70} The EGT model with parameters θ, f_θ , takes $x = (\mathcal{V}, \mathcal{E}, A)$ as an input and computes the embeddings for the atoms and bonds in the molecule. Specifically, for the atoms $\{v_1, v_2, \dots, v_N\}$ in x , we obtain the corresponding atom embeddings $\{h_1, h_2, \dots, h_N\}$, where each $h_l \in \mathbb{R}^{d_h}$ is an embedding for the l -th atom in the sample and d_h is the embedding dimension. Similarly, for bonds $e_{ij} \in \mathcal{E}$, we obtain the bond embeddings $H_{ij} \in \mathbb{R}^{d_H}$, where d_H is the dimension of the bond embeddings. All of these are sequentially updated through the EGT layers to produce the final embeddings $\{h_o^{[L]} \mid o = 1, \dots, N\}$ and $\{H_{ij}^{[L]} \mid \text{for } e_{ij} \in \mathcal{E}\}$, where L is the number of layers in the model. See Table S7 for a more detailed description of the data preprocessing and model architecture.

Training and inference

One straightforward way to train our deep graph neural network f_θ is to solve the regression problem of predicting knockdown efficacy. More formally, one can optimize the model parameter θ by minimizing the mean-squared error:

$$\min_{\theta} \frac{1}{m} \sum_{i=1}^m (f_{\theta}(x^{(i)}) - y^{(i)})^2 \quad (\text{Equation 3})$$

where $f_{\theta}(x^{(i)})$ and $y^{(i)}$ are the predicted and experimental knockdown efficacies, respectively. Note, however, that our database \mathcal{D} is a collection of multiple sources of experimental results and that inhibition rates differ by the environmental conditions, such as cell line, uptake, transfection, dosage, and cell density (cells/well), under which the experiment is conducted. Therefore, it is inappropriate to consider such optimization in our problem setting, where the input $x^{(i)}$ does not contain environmental factors but molecular features only.

Instead, we trained f_θ in a learning-to-rank fashion,⁷¹ a technique widely used in the ranking community.^{72,73} Simply put, learning to rank tries to learn a model f_θ that preserves the ordering in the original data: $f_{\theta}(x^{(i)}) > f_{\theta}(x^{(j)})$ if $y^{(i)} > y^{(j)}$. To this end, we fed the model two samples at a time and forced it to choose the more probable one. This comparison-based strategy makes more sense under a clever choice of a pair of samples. For example, we can randomly sample two instances that have the same experimental setups and nucleotide sequence, but with different chemical modifications. In this controlled scenario, a model can be trained to choose the more promising chemical modification, hiding the effects of experimental setups in the knockdown efficacy. We refer to the instance with higher regulatory efficacy as x_{pos} and the other as x_{neg} . Under this setting, we used the hinge loss on a pair of examples:

$$\mathcal{L}(\theta; x_{\text{pos}}, x_{\text{neg}}) = \max(0, t + f_{\theta}(x_{\text{neg}}) - f_{\theta}(x_{\text{pos}})) \quad (\text{Equation 4})$$

where $f_{\theta}(x_{\text{neg}}) - f_{\theta}(x_{\text{pos}})$ captures the pairwise difference in model predictions between the negative and positive samples. t is the threshold for the margin in the hinge loss. We obtained the θ^* that minimized the following loss function:

$$\theta^* = \operatorname{argmin}_{\theta} \sum_{i=1}^{m_{\text{pairs}}} \max(0, t + f_{\theta}(x_{\text{neg}}^{(i)}) - f_{\theta}(x_{\text{pos}}^{(i)})) \quad (\text{Equation 5})$$

We split the entire dataset \mathcal{D} into training and test sets at an 8:2 ratio based on the nucleotide sequences. Note that samples with a common nucleotide sequence, but different chemical modifications, were in the same set. To produce training pairs, we first chose one nucleotide sequence and then selected two samples from the same experimental setup, but with distinct chemical modifications, in which the difference in inhibition rates was larger than 10%. This process was repeated for all possible examples and pairs in the training set to yield $\mathcal{D}_{\text{pairs}}^{\text{train}}$. Similarly, we generated a test set $\mathcal{D}_{\text{pairs}}^{\text{test}}$. The number of pairs in $\mathcal{D}_{\text{pairs}}^{\text{train}}$ and $\mathcal{D}_{\text{pairs}}^{\text{test}}$ were 34,055 and 4,027, respectively. See Tables S8

and S9 for a statistical summary of the pair dataset used for chemical engineering.

During model training, we employed L2 regularization and early stopping to prevent overfitting. In detail, we used regularization strength 10^{-2} and we terminated the training iteration when the test accuracy plateaued over five training epochs. See Figure S8 for a sample learning curve where early stopping was applied.

Once our EGT-based model was trained using Equation 5, we utilized the model on a set of unseen examples to produce the rankings. Based on the rankings, we sorted out the most promising chemical modification candidates for further *in vitro* experiments.

Implementation details

The computational models were implemented in Python with TensorFlow 2.5.⁷⁴ The maximum number of atoms in our ASO database was $N = 516$. Our EGT model had $L = 12$ number of layers, and the embedding dimensions were $d_i = 64$ and $d_H = 64$. The model parameters were initialized under Xavier's rule⁷⁵ and updated with an RMSProp optimizer⁷⁶ during training, with a constant learning rate of 10^{-5} . All these hyper-parameters were found based on a grid search, which involved training the EGT-based model using all possible combinations of hyper-parameters considered. After three repeated training run for each hyper-parameter combination, we assessed the combination based on the hinge loss in Equation 5, on the test examples the corresponding model achieved; the best hyper-parameters were chosen based on the lowest test hinge loss. Types and respective ranges of hyper-parameters considered are summarized in Table S10. All experiments were conducted on a server with Intel Xeon Silver 4114 CPU @ 2.20 GHz and Nvidia TITAN RTX GPU.

Cell line and culture conditions

HeLa (cervical cancer) cells were purchased from the American Type Culture Collection (ATCC, USA) and maintained in DMEM (Welgene, Korea), which contains 10% FBS (Hyclone, USA) and 1% penicillin-streptomycin (P/S, 10,000 U/mL, Gibco, USA). THP-1 (acute monocytic leukemia) cells were obtained from Korean Cell Line Bank (KCLB, Korea) and maintained in RPMI-1640 (Welgene, Korea) that contains 10% FBS and 1% P/S. The cells were incubated in a humidified CO₂ incubator and subcultured every 3 days using 0.05% Trypsin-EDTA (Gibco, USA) once confluent.

Cell culture and THP-1 differentiation

HeLa cells and THP-1 cells were seeded in 24-well plates at a density of 3.5×10^4 and 2×10^5 cells/well, respectively. The THP-1 cells were treated with 100 nM 12-O-tetradecanoyl phorbol-13-acetate (PMA; Sigma-Aldrich, USA) for 72 h to induce differentiation into M0 macrophages and the medium was replaced with fresh medium after washing with PBS three times. To differentiate THP-1 cells into M1 macrophages, the PMA-treated THP-1 cells were incubated with 20 ng/mL rIFN- γ (Thermo Fisher, USA) and 10 ng/mL LPS (Sigma-Aldrich, USA) for 48 h.

WBs

HeLa cells, U87-MG cells, and THP-1 cells were transiently transfected with indicated concentrations of PS-ASO, gapmer, or ASOpt using lipofectamine RNAiMAX (Invitrogen, USA) based on the manufacturer's instructions. After incubating for 48 h, the cells were harvested with a scraper (SPL, USA), centrifuged at 1,500 rpm for 3 min, and lysed with radioimmunoprecipitation assay (RIPA) buffer (iNtRON, USA). The lysates were centrifuged at 13,000 rpm for 10 min and the supernatant was transferred to a new 1.5-mL tube. The protein levels were quantified with the BCA Protein Assay Kit (Thermo Scientific, USA). The protein samples were incubated in 5 \times SDS-PAGE loading buffer (Biosesang, Korea) at 95 $^\circ$ C for 10 min and separated on 10% Tris-glycine gels for 90 min at 100 V. The proteins were transferred onto Amersham Protran 0.45 NC membranes (Cytiva, USA) for 1.5 h at 300 mA on ice. The membranes were blocked in 5% Bovine Serum Albumin (GenDEPOT, USA) for 30 min at room temperature and washed three times with 1 \times Tris-buffered saline containing 0.1 % (w/v) Tween $^{\text{R}}$ 20 Detergent buffer (TBST; 20 mM Tris, 150 mM NaCl, pH 7.6) (MERK, USA). The membranes were incubated with primary IDO1 antibody (H-11, Santa Cruz Biotechnology, USA) for 2 h at room temperature and washed five times with 1 \times TBST. The membranes were then incubated with HRP-linked secondary antibody (Cell signaling, USA) for 1 h at room temperature and washed with 1 \times TBST five times. The images of bands were captured using the ImageQuant LAS-500 and an analysis of the normalized band density was performed with the ImageJ software.

qRT-PCR

HeLa and THP-1 cells were transiently transfected with the indicated concentrations of PS, gapmer, or ASOpt using lipofectamine RNAiMAX based on the manufacturer's instructions. After incubating for 48 h, the cells were harvested with a scraper, centrifuged at 1,500 rpm for 3 min, and lysed with TRIzol Reagent (Invitrogen, USA). The lysates were mixed with chloroform (Sigma-Aldrich, USA) and were centrifuged at 13,000 rpm for 10 min. The supernatant was transferred to a new 1.5-mL tube and mixed with 2-propanol EMSURE (Merk, USA). The mixtures were centrifuged at 13,000 rpm for 10 min and then washed with 70% ethanol. After removing the ethanol completely, the RNA pellet was dissolved in UltraPure DNase/RNase-Free distilled water (Invitrogen, USA). The concentration of the samples was measured using NanoDrop (Thermo Scientific, USA). The total RNA was synthesized into cDNA using a TOPscript RT DryMIX (Enzynomics, Korea). The expression levels of interesting genes were analyzed using TOPreal qRT-PCR 2 \times PreMIX (Enzynomics, Korea) on a CFX connect system (Bio-Rad, USA). The primers are listed in Table S11.

Cellular IDO1 enzyme assay

HeLa cells were seeded in 24-well plates at a density of 3×10^4 cells/well. The cells were stimulated with human rIFN- γ (50 ng/mL, R&D Systems, USA) and treated with PS-ASO, gapmer, and ASOpt, which were transfected by RNAiMAX at the indicated concentrations, or a negative control (only RNAiMAX) for 48 h in complete DMEM. Next, the supernatant (125 μ L) from each well was transferred to a

new 96-well plate, and 25 μL of trichloroacetic acid (30%) was added. The mixtures were heated at 65°C for 15 min, and the plates were centrifuged at 4,000 rpm for 15 min. The supernatant (100 μL) was carefully transferred to a new 96-well plate. p-Dimethylaminobenzaldehyde (2%, v/v) in acetic acid (100 μL) was added to the plates and the absorbance was read at a wavelength of 480 nm using a SpectraMAX-190 ELISA plate reader (Molecular Devices, USA).

Cytotoxicity assay

HeLa cells were transfected with the indicated concentrations of PS, gapmer, or ASOpt using lipofectamine RNAiMAX based on the manufacturer's instructions. Cytotoxicity was measured by the release of LDH using an LDH Assay Kit (Abcam, USA). Briefly, after incubating with the ASOs for 48 h, the 100 μL of medium from each well was harvested and centrifuged at 2,000 rpm for 3 min to remove debris. The centrifuged cells were carefully transferred to a new 96-well solid white plate (Corning, USA), and 100 μL LDH reaction mixture was added to each well. After incubating for 30 min, the absorbance was measured at 450 nm wavelength using a SpectraMAX-190 ELISA plate reader. Data analysis function: cytotoxicity (%) = (test sample – low control)/(high control – low control) \times 100; low control, non-treated sample; high control, cell lysis solution-treated sample.

Statistical analysis

Experimental results were analyzed using GraphPad Prism 10 software. All data are presented as the mean \pm standard deviation and computed from at least three independent trials. We consider our findings statistically significant if the *p* values were in the range $*p < 0.05$, $**p < 0.01$, using an unpaired two-tailed student's *t* test.

DATA AND CODE AVAILABILITY

Code is available at <https://github.com/Spidercores/ASOptimizer> (GitHub).

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.omtn.2024.102186>.

ACKNOWLEDGMENTS

This work was supported by the Technology Development Program (S3284154) funded by the Ministry of SMEs and Startups (MSS, Korea) and the Korea Bio Data Station(K-BDS) with computing resources, including technical support. We thank the members of the Biorchestra Corporation for their comments and suggestions. Additionally, we extend our thanks to Chen Fangying for her valuable support in the preparation of the manuscript. Figure 2A was created with BioRender.com.

AUTHOR CONTRIBUTIONS

G.H. conceived the project idea, developed ASOptimizer, analyzed data with statistical and deep-learning models, and mainly wrote the paper. M.C.K. managed the overall biological experiments and wrote the paper. D.S. conceived the project idea and edited the paper. D.H.K. conceived the project idea and supervised the work. D.L. built

the experimental database for ASOptimizer. K.L. conceived the project idea and supervised the work. E.K. conducted the *in vitro* experiments. M.G.K. conceived the project idea, developed ASOptimizer, supervised the work, analyzed data with statistical and deep-learning models, and wrote the paper. J.-H.R. conceived the project idea and supervised the work.

DECLARATION OF INTERESTS

G.H., D.S., D.L., K.L., and M.G.K. are employees of Spidercore Inc., and M.C.K., D.H.K., E.K., and J.-H.R. are employees of BIORCHESTRA Co., Ltd.

REFERENCES

- Walder, R.Y., and Walder, J.A. (1988). Role of RNase H in hybrid-arrested translation by antisense oligonucleotides. *Proc. Natl. Acad. Sci. USA* 85, 5011–5015.
- Uhlmann, E., and Peyman, A. (1990). Antisense oligonucleotides: a new therapeutic principle. *Chem. Rev.* 90, 543–584.
- Dias, N., and Stein, C.A. (2002). Antisense oligonucleotides: basic concepts and mechanisms. *Mol. Cancer Therapeut.* 1, 347–355.
- Parham, J.S., and Goldberg, A.C. (2019). Mipomersen and its use in familial hypercholesterolemia. *Expert Opin. Pharmacother.* 20, 127–131.
- Khan, N., Eliopoulos, H., Han, L., Kinane, T.B., Lowes, L.P., Mendell, J.R., Gordish-Dressman, H., Henricson, E.K., and McDonald, C.M.; Eteplirsen Investigators and the CINRG DNHS Investigators and the CINRG DNHS Investigators (2019). Eteplirsen Treatment Attenuates Respiratory Decline in Ambulatory and Non-Ambulatory Patients with Duchenne Muscular Dystrophy. *J. Neuromuscul. Dis.* 6, 213–225.
- Neil, E.E., and Bisaccia, E.K. (2019). Nusinersen: A Novel Antisense Oligonucleotide for the Treatment of Spinal Muscular Atrophy. *J. Pediatr. Pharmacol. Therapeut.* 24, 194–203.
- Mathew, V., and Wang, A.K. (2019). Inotersen: new promise for the treatment of hereditary transthyretin amyloidosis. *Drug Des. Dev. Ther.* 13, 1515–1525.
- Scott, L.J. (2020). Givosiran: First Approval. *Drugs* 80, 335–339.
- Crooke, S.T. (2004). Progress in antisense technology. *Annu. Rev. Med.* 55, 61–95.
- Chan, J.H.P., Lim, S., and Wong, W.S.F. (2006). Antisense oligonucleotides: from design to therapeutic application. *Clin. Exp. Pharmacol. Physiol.* 33, 533–540.
- Kole, R., Krainer, A.R., and Altman, S. (2012). RNA therapeutics: beyond RNA interference and antisense oligonucleotides. *Nat. Rev. Drug Discov.* 11, 125–140.
- Hagedorn, P.H., Persson, R., Funder, E.D., Albæk, N., Diemer, S.L., Hansen, D.J., Møller, M.R., Papargyri, N., Christiansen, H., Hansen, B.R., et al. (2018). Locked nucleic acid: modality, diversity, and drug discovery. *Drug Discov. Today* 23, 101–114.
- Papargyri, N., Pontoppidan, M., Andersen, M.R., Koch, T., and Hagedorn, P.H. (2020). Chemical Diversity of Locked Nucleic Acid-Modified Antisense Oligonucleotides Allows Optimization of Pharmaceutical Properties. *Mol. Ther. Nucleic Acids* 19, 706–717.
- Crooke, S.T., Baker, B.F., Crooke, R.M., and Liang, X.H. (2021). Antisense technology: an overview and prospectus. *Nat. Rev. Drug Discov.* 20, 427–453.
- Yoshida, T., Morihiro, K., Naito, Y., Mikami, A., Kasahara, Y., Inoue, T., and Obika, S. (2022). Identification of nucleobase chemical modifications that reduce the hepatotoxicity of gapmer antisense oligonucleotides. *Nucleic Acids Res.* 50, 7224–7234.
- Yoshida, T., Naito, Y., Yasuhara, H., Sasaki, K., Kawaji, H., Kawai, J., Naito, M., Okuda, H., Obika, S., and Inoue, T. (2019). Evaluation of off-target effects of gapmer antisense oligonucleotides using human cells. *Gene Cell.* 24, 827–835.
- Elbashir, S.M., Harborth, J., Weber, K., and Tuschl, T. (2002). Analysis of gene function in somatic mammalian cells using small interfering RNAs. *Methods* 26, 199–213.
- Amarzguoui, M., and Prydz, H. (2004). An algorithm for selection of functional siRNA sequences. *Biochem. Biophys. Res. Commun.* 316, 1050–1058.
- Reynolds, A., Leake, D., Boese, Q., Scaringe, S., Marshall, W.S., and Khvorova, A. (2004). Rational siRNA design for RNA interference. *Nat. Biotechnol.* 22, 326–330.

20. Ui-Tei, K., Naito, Y., Takahashi, F., Haraguchi, T., Ohki-Hamazaki, H., Juni, A., Ueda, R., and Saigo, K. (2004). Guidelines for the selection of highly effective siRNA sequences for mammalian and chick RNA interference. *Nucleic Acids Res.* 32, 936–948.
21. Naito, Y., and Ui-Tei, K. (2012). siRNA Design Software for a Target Gene-Specific RNA Interference. *Front. Genet.* 3, 102.
22. Sciabola, S., Cao, Q., Orozco, M., Faustino, I., and Stanton, R.V. (2013). Improved nucleic acid descriptors for siRNA efficacy prediction. *Nucleic Acids Res.* 41, 1383–1394.
23. Saetrom, P. (2004). Predicting the efficacy of short oligonucleotides in antisense and RNAi experiments with boosted genetic programming. *Bioinformatics* 20, 3055–3063.
24. Cui, W., Ning, J., Naik, U.P., and Duncan, M.K. (2004). OptiRNAi, an RNAi design tool. *Comput. Methods Progr. Biomed.* 75, 67–73.
25. Teramoto, R., Aoki, M., Kimura, T., and Kanaoka, M. (2005). Prediction of siRNA functionality using generalized string kernel and support vector machine. *FEBS Lett.* 579, 2878–2882.
26. Huesken, D., Lange, J., Mickanin, C., Weiler, J., Asselbergs, F., Warner, J., Meloan, B., Engel, S., Rosenberg, A., Cohen, D., et al. (2005). Design of a genome-wide siRNA library using an artificial neural network. *Nat. Biotechnol.* 23, 995–1001.
27. Shabalina, S.A., Spiridonov, A.N., and Ogurtsov, A.Y. (2006). Computational models with thermodynamic and composition features improve siRNA design. *BMC Bioinf.* 7, 65.
28. Matveeva, O., Nechipurenko, Y., Rossi, L., Moore, B., Saetrom, P., Ogurtsov, A.Y., Atkins, J.F., and Shabalina, S.A. (2007). Comparison of approaches for rational siRNA design leading to a new efficient and transparent method. *Nucleic Acids Res.* 35, e63.
29. Harding, P.L., Fall, A.M., Honeyman, K., Fletcher, S., and Wilton, S.D. (2007). The influence of antisense oligonucleotide length on dystrophin exon skipping. *Mol. Ther.* 15, 157–166.
30. Popplewell, L.J., Trollet, C., Dickson, G., and Graham, I.R. (2009). Design of phosphorodiamidate morpholino oligomers (PMOs) for the induction of exon skipping of the human DMD gene. *Mol. Ther.* 17, 554–561.
31. Echigoya, Y., Mouly, V., Garcia, L., Yokota, T., and Duddy, W. (2015). In silico screening based on predictive algorithms as a design tool for exon skipping oligonucleotides in Duchenne muscular dystrophy. *PLoS One* 10, e0120058.
32. Chiba, S., Lim, K.R.Q., Sheri, N., Anwar, S., Erkut, E., Shah, M.N.A., Aslesh, T., Woo, S., Sheikh, O., Maruyama, R., et al. (2021). eSkip-Finder: a machine learning-based web application and database to identify the optimal sequences of antisense oligonucleotides for exon skipping. *Nucleic Acids Res.* 49, W193–W198.
33. Ding, Y., Chan, C.Y., and Lawrence, C.E. (2004). Sfold web server for statistical folding and rational design of nucleic acids. *Nucleic Acids Res.* 32, W135–W141.
34. Shao, Y., Wu, Y., Chan, C.Y., McDonough, K., and Ding, Y. (2006). Rational design and rapid screening of antisense oligonucleotides for prokaryotic gene modulation. *Nucleic Acids Res.* 34, 5660–5669.
35. Sciabola, S., Xi, H., Cruz, D., Cao, Q., Lawrence, C., Zhang, T., Rotstein, S., Hughes, J.D., Caffrey, D.R., and Stanton, R.V. (2021). PFRED: A computational platform for siRNA and antisense oligonucleotides design. *PLoS One* 16, e0238753.
36. <http://www.genscript.com/tools/sirna-target-finder>.
37. <https://rnaidesigner.thermofisher.com/rnaexpress/>.
38. Enright, A.J., John, B., Gaul, U., Tuschl, T., Sander, C., and Marks, D.S. (2003). MicroRNA targets in *Drosophila*. *Genome Biol.* 5, R1.
39. Lewis, B.P., Shih, I.H., Jones-Rhoades, M.W., Bartel, D.P., and Burge, C.B. (2003). Prediction of mammalian microRNA targets. *Cell* 115, 787–798.
40. Rehmsmeier, M., Steffen, P., Hochsmann, M., and Giegerich, R. (2004). Fast and effective prediction of microRNA/target duplexes. *RNA* 10, 1507–1517.
41. Krek, A., Grün, D., Poy, M.N., Wolf, R., Rosenberg, L., Epstein, E.J., MacMenamin, P., da Piedade, I., Gunsalus, K.C., Stoffel, M., and Rajewsky, N. (2005). Combinatorial microRNA target predictions. *Nat. Genet.* 37, 495–500.
42. Miranda, K.C., Huynh, T., Tay, Y., Ang, Y.S., Tam, W.L., Thomson, A.M., Lim, B., and Rigoutsos, I. (2006). A pattern-based method for the identification of MicroRNA binding sites and their corresponding heteroduplexes. *Cell* 126, 1203–1217.
43. Opitz, C.A., Somarrivas Patterson, L.F., Mohapatra, S.R., Dewi, D.L., Sadik, A., Platten, M., and Trump, S. (2020). The therapeutic potential of targeting tryptophan catabolism in cancer. *Br. J. Cancer* 122, 30–44.
44. Peng, X., Zhao, Z., Liu, L., Bai, L., Tong, R., Yang, H., and Zhong, L. (2022). Targeting Indoleamine Dioxigenase and Tryptophan Dioxigenase in Cancer Immunotherapy: Clinical Progress and Challenges. *Drug Des. Dev. Ther.* 16, 2639–2657.
45. Ye, J., Kumanova, M., Hart, L.S., Sloane, K., Zhang, H., De Panis, D.N., Bobrovnikova-Marjon, E., Diehl, J.A., Ron, D., and Koumenis, C. (2010). The GCN2-ATF4 pathway is critical for tumour cell survival and proliferation in response to nutrient deprivation. *EMBO J.* 29, 2082–2096.
46. Tomek, P., Gore, S.K., Potts, C.L., Print, C.G., Black, M.A., Hallermayr, A., Kilian, M., Sattlegger, E., and Ching, L.M. (2018). Imprinted and ancient gene: a potential mediator of cancer cell survival during tryptophan deprivation. *Cell Commun. Signal.* 16, 88.
47. Esser, C. (2016). The Aryl Hydrocarbon Receptor in Immunity: Tools and Potential. *Methods Mol. Biol.* 1371, 239–257.
48. Xue, P., Fu, J., and Zhou, Y. (2018). The Aryl Hydrocarbon Receptor and Tumor Immunity. *Front. Immunol.* 9, 286.
49. <https://www.jens.org/>.
50. RICHARD, K. and FRANK, J. Immunosuppression-Reverting Oligonucleotides Inhibiting the Expression of IDO.
51. Yasuhara, H., Yoshida, T., Sasaki, K., Obika, S., and Inoue, T. (2022). Reduction of Off-Target Effects of Gapmer Antisense Oligonucleotides by Oligonucleotide Extension. *Mol. Diagn. Ther.* 26, 117–127.
52. Sugimoto, N., Nakano, S., Katoh, M., Matsumura, A., Nakamuta, H., Ohmichi, T., Yoneyama, M., and Sasaki, M. (1995). Thermodynamic parameters to predict stability of RNA/DNA hybrid duplexes. *Biochemistry* 34, 11211–11216.
53. Hussain, M.S., Zaki, M.J., and Subramanian, D. (2022). Global self-attention as a replacement for graph convolution. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 655–665.
54. Italiani, P., and Boraschi, D. (2014). From Monocytes to M1/M2 Macrophages: Phenotypical vs. Functional Differentiation. *Front. Immunol.* 5, 514.
55. Liu, J., Geng, X., Hou, J., and Wu, G. (2021). New insights into M1/M2 macrophages: key modulators in cancer progression. *Cancer Cell Int.* 21, 389.
56. Wang, X.F., Wang, H.S., Wang, H., Zhang, F., Wang, K.F., Guo, Q., Zhang, G., Cai, S.H., and Du, J. (2014). The role of indoleamine 2,3-dioxygenase (IDO) in immune tolerance: focus on macrophage polarization of THP-1 cells. *Cell. Immunol.* 289, 42–48.
57. Duan, Z., and Luo, Y. (2021). Targeting macrophages in cancer immunotherapy. *Signal Transduct. Targeted Ther.* 6, 127.
58. Rouillard, J.M., Zuker, M., and Gulari, E. (2003). OligoArray 2.0: design of oligonucleotide probes for DNA microarrays using a thermodynamic approach. *Nucleic Acids Res.* 31, 3057–3062.
59. Chen, X., Li, Y., Umarov, R., Gao, X., and Song, L. (2020). RNA Secondary Structure Prediction By Learning Unrolled Algorithms. In *International Conference on Learning Representations*.
60. Fu, L., Cao, Y., Wu, J., Peng, Q., Nie, Q., and Xie, X. (2022). UFold: fast and accurate RNA secondary structure prediction with deep learning. *Nucleic Acids Res.* 50, e14.
61. Sato, K., Akiyama, M., and Sakakibara, Y. (2021). RNA secondary structure prediction using deep learning with thermodynamic integration. *Nat. Commun.* 12, 941.
62. Clever, J., Sasseti, C., and Parslow, T.G. (1995). RNA secondary structure and binding sites for gag gene products in the 5' packaging signal of human immunodeficiency virus type 1. *J. Virol.* 69, 2101–2109.
63. O'Leary, N.A., Wright, M.W., Brister, J.R., Ciufu, S., Haddad, D., McVeigh, R., Rajput, B., Robbertse, B., Smith-White, B., Ako-Adjei, D., et al. (2016). Reference sequence (RefSeq) database at NCBI: current status, taxonomic expansion, and functional annotation. *Nucleic Acids Res.* 44, D733–D745.
64. Zuker, M., Mathews, D.H., and Turner, D.H. (1999). Algorithms and Thermodynamics for RNA Secondary Structure Prediction: A Practical Guide (Springer Netherlands), pp. 11–43.

65. Mills, N. (2006). ChemDraw Ultra 10.0 CambridgeSoft.
66. Landrum, G. (2023). RDKit. GitHub Repository Accessed.
67. Hu, W., Fey, M., Zitnik, M., Dong, Y., Ren, H., Liu, B., Catasta, M., and Leskovec, J. (2020). Open Graph Benchmark: Datasets for Machine Learning on Graphs. *Adv. Neural Inf. Process. Syst.* 33, 22118–22133.
68. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L.u., and Polosukhin, I. (2017). Attention is All you Need. In *Advances in Neural Information Processing Systems*.
69. Nakata, M., and Shimazaki, T. (2017). PubChemQC project: a large-scale first-principles electronic structure database for data-driven chemistry. *J. Chem. Inf. Model.* 57, 1300–1308.
70. Dwivedi, V.P., Joshi, C.K., Laurent, T., Bengio, Y., and Bresson, X. (2023). Benchmarking graph neural networks. *J. Mach. Learn. Res.* 24, 1–48.
71. Liu, T.-Y. (2007). Learning to rank for information retrieval. *FNT. in Information Retrieval* 3, 225–331.
72. Li, Y., Song, Y., and Luo, J. (2017). Improving pairwise ranking for multi-label image classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3617–3625.
73. Doughty, H., Damen, D., and Mayol-Cuevas, W. (2018). Who's better? who's best? pairwise deep ranking for skill determination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6057–6066.
74. Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al. (2016). TensorFlow: A System for Large-Scale Machine Learning. In *Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation*, pp. 265–283.
75. Glorot, X., and Bengio, Y. (2010). Understanding the difficulty of training deep feed-forward neural networks. In *Proceedings of the 13th International Conference on Artificial Intelligence and Statistics*, pp. 249–256.
76. Tieleman, T., and Hinton, G. (2012). Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural networks for machine learning* 4, 26–31.