004

Discovering Hidden Algebraic Structures via Transformers with Rank-Aware Beam GRPO

Anonymous Authors¹

Abstract

Recent efforts have extended the capabilities of transformers in logical reasoning and symbolic computations. In this work, we investigate their capacity for non-linear latent pattern discovery 015 in the context of functional decomposition, fo-016 cusing on the challenging algebraic task of mul-018 tivariate polynomial decomposition. This problem, with widespread applications in science and engineering, is proved to be NP-hard, and demands both precision and insight. Our contributions are threefold: First, we develop a synthetic data generation pipeline providing fine-grained control over problem complexity. Second, we 024 train transformer models via supervised learning 025 and evaluate them across four key dimensions involving scaling behavior and generalizability. 028 Third, we propose Beam Grouped Relative Policy 029 Optimization (BGRPO), a rank-aware reinforcement learning method suitable for hard algebraic problems. Finetuning with BGRPO improves accuracy while reducing beam width by up to half, resulting in approximately 75% lower inference 034 compute. Additionally, our model demonstrates competitive performance in polynomial simplification, outperforming Mathematica in various cases.

1. Introduction

038

043

044

045

047

048

051

053

Transformers, initially developed for natural language processing (Vaswani et al., 2017), have shown remarkable versatility across diverse domains such as vision (Dosovitskiy et al., 2020) and protein folding (Jumper et al., 2021). More recently, their applications in formal reasoning, symbolic mathematics and algorithmic tasks start to gain traction. Several works have showcased transformer-based architectures'

ability to tackle highly structured problems, including theorem proving (Polu & Sutskever, 2020; Trinh et al., 2024), integration (Lample & Charton, 2020), matrix multiplication (Fawzi et al., 2022) and equation solving (Drori et al., 2022).

In this work, we investigate the transformer's capacity for non-linear latent pattern discovery in the context of functional decomposition, i.e. decomposing a complex function as the composition of simpler sub-functions. In contrast to step-by-step logical deduction, or pattern recognition in data analysis, functional decomposition poses significant new challenges to the transformer, because the forms of the sub-functions that we try to discover can be totally hidden or obscured in the final compact form of the original function. Furthermore, it requires extreme precision without any margin of error. Unlike more forgiving classification tasks, the decomposition problem admits only a sparse set of correct solutions: even minor deviations in signs or coefficients can render outputs completely invalid.

Beyond its theoretical interest, functional decomposition has ubiquitous applications in software engineering (Tempero et al., 2024), systems biology (Mori et al., 2023), mechanical design (She et al., 2024), systems engineering (Hernandez et al., 2024) and digital logic design (Adamski et al., 2005; Lin et al., 2008), where capturing hidden substructures within high-dimensional functions leads to more tractable and efficient models. However, identifying a function's latent compositional structure requires models to look past surface-level correlations, attending instead to deep algebraic symmetries and invariants.

A particularly rich case of functional decomposition arises in multivariate polynomial functions. The polynomial decomposition problem over a ring k seeks to decompose a given polynomial $f \in k[x_1, \ldots, x_n]$ into polynomials $g \in k[y_1, \ldots, y_m]$ and $h_1, \ldots, h_m \in k[x_1, \ldots, x_n]$ such that

$$f(x_1, \dots, x_n) = g(h_1(x_1, \dots, x_n), \dots, h_m(x_1, \dots, x_n)).$$
(1)

It has wide-ranging applications from cryptography (Patarin & Goubin, 1997) to dynamical modeling (Dang & Testylier, 2012), signal processing (Demirtas et al., 2012) and robotics (Elias & Wen, 2025; Manocha & Canny, 1992).

The multivariate polynomial decomposition problem has

¹Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

been proved to be NP-hard by Dickerson (Dickerson, 1987;
1993), although efficient algorithms for various special cases
are discussed in (Gathen et al., 2003; Von Zur Gathen,
1990a;b; Faugère & Perret, 2009a;b; Zhao et al., 2012).
To illustrate the difficulty of the problem for the models, let
us consider the following expression

$$\begin{split} f &= 2a_1^3b_1^3 + 25a_1^2b_1^2 + 6a_1^2a_2b_2b_1^2 + 6a_1^2a_3b_3b_1^2 + 6a_1a_2^2b_2^2b_1 \\ &\quad + 6a_1a_3^2b_3^2b_1 + 96a_1b_1 + 50a_1a_2b_2b_1 + 50a_1a_3b_3b_1 \\ &\quad + 12a_1a_2a_3b_2b_3b_1 + 2a_2^3b_2^3 + 2a_3^3b_3^3 + 25a_2^2b_2^2 + 25a_3^2b_3^2 \\ &\quad + 6a_2a_3^2b_2b_3^2 + 96a_2b_2 + 6a_2^2a_3b_2^2b_3 + 96a_3b_3 \\ &\quad + 50a_2a_3b_2b_3 + 12 \end{split}$$

061

062 063

064

065

066

067 It has a hidden O(3)-symmetry, which can be revealed by 068 decomposing $f = g \circ h$, with $g(y) = y^2 + 2(4+y)^3$ and 069 $h = a_1b_1 + a_2b_2 + a_3b_3$. This is a highly nontrivial task 070 to identify the inner function h directly from the expanded 071 form of f, as its structure becomes completely obscured after polynomial substitution, expansion and simplification. Even in this relatively constrained case where q is univari-074 ate, discovering the decomposition requires recognizing 075 non-linear latent patterns across dozens of terms. When 076 g becomes multivariate, the complexity increases substan-077 tially, making the problem even more challenging. 078

To tackle the polynomial decomposition problem, we de-079 velop a systematic approach with four key components. First, we create a backward synthetic data generation 081 pipeline that allows fine-grained control over polynomial 082 complexity involving range of coefficients, degree, and num-083 ber of variables. Second, we train lightweight transformer models on these synthetic datasets using supervised learn-085 ing and analyze how performance scales across four axes (performance complexity scaling, architecture scaling, dis-087 tribution adaptation, search strategy analysis). Third, we discover that both multi-sampling and greedy search meth-089 ods struggle with the sparse solution space of the polyno-090 mial decomposition problem, and we implement a beam 091 search strategy to effectively extract the models' capabili-092 ties. Finally, to address the computational intensity of beam 093 search, we develop a rank-aware variant of the Grouped Rel-094 ative Policy Optimization (GRPO) reinforcement learning 095 algorithm, which encodes rank information directly in the 096 reward function. 097

098 Our study makes the following contributions to neural ap-099 proaches for polynomial decomposition. First, our back-100 ward data generation pipeline enables targeted training across varying levels of decomposition difficulty. Second, our comprehensive evaluation across four dimensions, for the first time, establishes robust baselines for transformers' 104 performance on polynomial decomposition tasks. Third, 105 using the rank-aware Beam Grouped Relative Policy Op-106 timization (BGRPO), our models improve accuracy while 107 reducing beam search width by up to 50%, resulting in 75% lower computational requirements during inference. Addi-109

tionally, our model demonstrates competitive performance in polynomial simplification, outperforming Mathematica in various cases. This underscores the potential of neural models to complement and extend classical symbolic computation capabilities.

2. Method

2.1. Backward Synthetic Data Generation

We generate synthetic data for supervised learning using a backward approach, starting from the decomposed form. First, we generate the inner functions (h_1, \ldots, h_m) in Eq. (1)) and the outer function (g in Eq. (1)) with random monomial terms of bounded degree and random coefficients within a given range. Then, we obtain the composed function (f in Eq. (1)) via substitution, expansion, and term collection. See Appendix A for the detailed algorithm. For each generated instance, we create a training pair consisting of the expanded polynomial f as input and its decomposed components $\{g, h_1, \ldots, h_{v_{outer}}\}$ as the target output. The model is trained to minimize the standard negative loglikelihood loss function.

Our synthetic data generation process provides fine-grained control over problem complexity through eight parameters: C_{inner} (coefficient range for inner polynomials), d_{inner} (maximum degree of inner polynomials), v_{inner} (number of variables in inner polynomials), t_{inner} (maximum number of terms in inner polynomials), and similarly C_{outer} , d_{outer} , v_{outer} , and t_{outer} for the outer polynomial.

2.2. Beam Search

Beam search is a breadth-first search algorithm that approximates optimal decoding by keeping track of the k most probable sequences at each step (Freitag & Al-Onaizan, 2017). For each of the k current sequences, the algorithm considers the top-k token extensions per sequence. These k^2 candidate continuations are then ranked by the sum of log probabilities of all tokens in the sequence, and only the top-k sequences with the highest cumulative log probability are retained for the next step. In this paper, we refer to kas the beam width, and to the position (1st, 2nd, etc.) of an output in the final beam as its rank.

Our analysis across all model outputs identified a specific error pattern in polynomial decomposition: the model achieves approximately 90% accuracy for predicting nonsign tokens (operators, numbers, variables), but exhibits near-random performance for deciding between positive and negative signs. This creates a unique inference challenge where exploration needs to be constrained for highconfidence structural elements while simultaneously expanded for uncertain sign choices.

Beam search is particularly well-suited for this situation as 111 it maintains the high-confidence structural backbone while 112 systematically exploring variations in the uncertain com-113 ponents. Our experiments demonstrate that beam search 114 significantly outperforms greedy decoding and random sampling for polynomial decomposition tasks. See Appendix C 115 116 for a detailed error analysis and an explanation of beam 117 search effectiveness for this task. 118

2.3. BGRPO : Reinforcement Learning Method 120 **Enhancing Beam Search Efficiency**

119

121 The computational cost of beam search scales quadratically 122 with beam width. There would be a significant compu-123 tational advantage if we could improve the ranks of cor-124 rect outputs. To address this, we introduce Beam Grouped 125 Relative Policy Optimization (BGRPO), a reinforcement 126 learning method that extends GRPO, uniquely taking into 127 account rankings in the beam search, specifically designed 128 for improving beam search inference efficiency. 129

130 Reinforcement learning enables models to explore solu-131 tion spaces more effectively than supervised learning alone, 132 enhancing the model's capabilities by addressing specific 133 weaknesses through a reward mechanism. This approach en-134 courages correct answers while discouraging incorrect ones 135 based on an advantage function-the difference between a 136 solution's reward and a baseline reward. Group Relative Pol-137 icy Optimization (GRPO) (Shao et al., 2024) estimates this 138 baseline for each question by sampling a group of outputs, 139 and has shown promising results for reinforcement learning 140 in language generation tasks due to its sample efficiency 141 and stability (DeepSeek-AI, 2025). We chose GRPO over 142 traditional Proximal Policy Optimization (PPO) (Schulman 143 et al., 2017) because it eliminates the need for a separate 144 value network or reward model, reducing training complex-145 ity while improving stability, and its group-wise baseline 146 calculation naturally fits tasks with a clear binary reward 147 structure like polynomial decomposition. 148

Our proposed Beam Grouped Relative Policy Optimiza-149 tion (BGRPO) extends this approach by using beam search 150 rather than independent sampling for generating the group 151 of outputs. While this significantly alters the distribution 152 of outputs, making their average reward less suitable as a 153 traditional baseline, it still provides valid training signals by 154 reinforcing correct answers and penalizing incorrect ones. 155 BGRPO is particularly effective for our task because beam 156 search generates outputs with identical structure that differ 157 only in the confusing elements (signs), creating a focused 158 learning signal. 159

160 Additionally, BGRPO incorporates rank information di-161 rectly into the reward function by applying an exponential 162 decay factor based on the position in the beam. This incen-163 tivizes correct answers to appear at earlier positions in the 164

beam search, effectively pushing correct solutions toward the top of the beam ranking.

Training Objective For a prompt x, let $\mathcal{B}(x) =$ $\{y_1, \ldots, y_w\}$ be the set of beam search outputs with beam width w generated by the old policy $\pi_{\theta_{\text{old}}}$. Each output sequence y_i receives a reward r_i , where $r_i = 0$ for incorrect polynomial decomposition and $r_i = 1$ for correct decomposition. In BGRPO, we incorporate rank information by scaling the reward for correct decompositions using an exponential decay function $e^{-\operatorname{rank}/w}$. We optimize the policy model π_{θ} for μ iterations by maximizing the following objective, $\mathcal{J}_{\text{BGRPO}}(\theta)$, where we define $\rho_i = \frac{\pi_{\theta}(y_i|x)}{\pi_{\theta_{\text{old}}}(y_i|x)}$ for brevity:

$$\mathcal{J}_{\text{BGRPO}}(\theta) = \frac{1}{w} \sum_{i=1}^{w} \left[\min\left(\rho_i A_i, \text{clip}(\rho_i, 1-\varepsilon, 1+\varepsilon) A_i\right) - \beta \mathbb{D}_{\text{KL}}(\pi_{\theta} || \pi_{\text{ref}}) \right]$$
(2)

where ε is the clipping parameter that constrains policy updates and β controls the KL divergence regularization term:

$$\mathbb{D}_{\mathrm{KL}}(\pi_{\theta}||\pi_{\mathrm{ref}}) = \frac{\pi_{\mathrm{ref}}(o_i|q)}{\pi_{\theta}(o_i|q)} - \log\frac{\pi_{\mathrm{ref}}(o_i|q)}{\pi_{\theta}(o_i|q)} - 1.$$
(3)

where ε is the clipping parameter that constrains policy updates and β controls the KL divergence regularization term:

$$\mathbb{D}_{\mathrm{KL}}(\pi_{\theta}||\pi_{\mathrm{ref}}) = \frac{\pi_{\mathrm{ref}}(o_i|q)}{\pi_{\theta}(o_i|q)} - \log\frac{\pi_{\mathrm{ref}}(o_i|q)}{\pi_{\theta}(o_i|q)} - 1.$$
(4)

Here, π_{ref} is the reference policy, which is the initial model before BGRPO training. The advantage function A_i is computed without normalization as $A_i = r_i - r_i$ $mean(\{r_1, r_2, \cdots, r_w\})$, following the approach in (Liu et al., 2025).

3. Experimental Setup

3.1. Evaluation Axes

To systematically analyze our models' capabilities for the polynomial decomposition problem, we consider four key evaluation dimensions.

Problem Complexity Scaling (\mathcal{D}_1). We analyze how the model performance varies with respect to changes in the complexity parameters for synthetic data generation. We vary the number of variables $v_{\mathrm{inner}}, v_{\mathrm{outer}}$, and the maximum degrees $d_{\text{inner}}, d_{\text{outer}}$ for both the inner and outer polynomials.

Architecture Scaling (D_2). We investigate how model performance scales with key architectural hyperparameters 165 of the transformer. In particular, we measure $\mathcal{P}(M(d, l, a))$, 166 the performance of models with embedding dimension d, 167 number of layers l, and number of attention heads a. Our 168 goal is to characterize how these hyperparameters influence 169 model capabilities.

170 **Distribution Adaptation** (\mathcal{D}_3). A practical challenge in 171 applying transformers to symbolic computation is their sen-172 sitivity to the numerical ranges present in the training data. 173 For example, models trained on specific coefficient ranges 174 tend to struggle with polynomials outside these ranges. On 175 the other hand, we found that models can rapidly adapt 176 to new coefficient distributions with minimal additional 177 training, suggesting that they manage to learn generalizable 178 pattern recognition rather than merely memorizing specific 179 numerical relationships. 180

181 To quantify the model's ability to transfer its polynomial 182 decomposition skills to numerically distinct but structurally 183 identical problems, we prepare the model $M_{C_1 \rightarrow C_2}^n$. This 184 model is initially trained on 1M polynomial decomposition examples with $C_{\text{outer}} = C_1$ and then fine-tuned with n185 186 examples with $C_{\text{outer}} = C_2$ where $C_1 \cap C_2 = \emptyset$. We 187 measure the performance of model $M^n_{C_1 \rightarrow C_2}$ on a test set of 188 polynomial decomposition problems with $C_{outer} = C_2$: 189

$$\mathcal{G}(n) = \mathcal{P}(M^n_{C_1 \to C_2}, \text{ test set with } C_{\text{outer}} = C_2)$$
 (5)

192 Search Strategy Analysis (\mathcal{D}_4). We investigate how beam 193 search enhances model performance on polynomial decom-194 position tasks, analyzing its effectiveness across different 195 model architectures and levels of problem complexity.

1973.2. Synthetic Dataset Setup

190

191

For the axis D_1 of the problem complexity scaling, we first examine degree scaling by training a model on 2M polynomial decomposition examples with different inner and outer degrees as described in Table 1. We then evaluate this model on separate test datasets with the same configuration parameters, each corresponding to one of nine different (d_{inner}, d_{outer}) pairs to assess performance across varying problem complexities.

For the second part of the D_1 axis, we train a model for each combination of v_{inner} and v_{outer} varying from 2 to 4 while fixing the other parameter at 3. For each combination, we use 1M examples to train the model.

212 For the axis D_2 of architecture scaling, we train multiple 213 models with varying architectural configurations, all using 214 the same dataset of 2M examples with polynomial parame-215 ters as described in Table 1.

For the axis \mathcal{D}_3 of distribution adaptation, we train initial models on 1M examples with $C_{\text{outer}} = C_1 = [-5,5]$ and then adapt them to examples with $C_{\text{outer}} = C_2 =$ $[-10, -6] \cup [6, 10]$. Other parameters are the same across both datasets as described in Table 1.

For the second part of \mathcal{D}_1 (Variable Scaling) and \mathcal{D}_2 , we set $t_{\text{inner}} = t_{\text{outer}} = 3$ to prevent expressions from becoming too long. We describe our tokenization in Appendix B.

Table 1. Synthetic Dataset Configuration Across Evaluation Axes

Evaluation Axis	Inner Coeff.	Outer Coeff.	Inner Degrees	Outer Degrees	Inner Vars	Outer Vars
D_1 (Degree) D_1 (Variable)	[-20, 20] [-5, 5]	[-20, 20] [-5, 5]	${2, 3, 4} \\ 3$	${2, 3, 4}$ 3	1 {2, 3, 4}	1 {2, 3, 4}
D_2 (Arch.)	[-5, 5]	[-5, 5]	3	3	3	3
\mathcal{D}_3	[-20, 20] [-20, 20]	$\begin{array}{c} C_1 = [-5,5] \\ C_2 = [-10,-6] \cup [6,10] \end{array}$	$\{1, 2\}$ $\{1, 2\}$	$\{1, 2, 3, 4\}$ $\{1, 2, 3, 4\}$	1	1

3.3. Architecture Configuration

We employ a decoder-only transformer architecture following standard design principles (Vaswani et al., 2017). Table 2 summarizes our task-specific configurations across all experimental axes. For lightweight and effective training, we developed our own model and training pipeline based on minGPT (Karpathy, 2020).

3.4. Supervised Learning Details

We train our models using the Adam optimizer with an initial learning rate of 6×10^{-4} , incorporating a 10% warmup period followed by cosine decay. Each configuration initially trains on 1M instances, with additional 1M training examples added incrementally until performance saturation. We use a batch size of 200 throughout training. We train models with enough epochs until it saturates with the given dataset.

3.5. BGRPO Implementation

For the BGRPO reinforcement learning phase, we generate candidate solutions using beam search with a width of 32 and temperature of 1.0. We implement our approach using the GRPO functionality from the trl library (von Werra et al., 2020). The training process consists of 5 policy update iterations after sampling outputs for 8 distinct polynomial decomposition problems. We set the PPO clipping parameter ε to 0.2 and the KL divergence coefficient β to 0.01. The learning rate during BGRPO training is 1×10^{-5} . We train models from \mathcal{D}_2 on a dataset of 200 non-repeating problems, saving checkpoints every 5 iterations and selecting the best model based on performance with beam width 7.

4. Experimental Results

4.1. Problem Complexity Scaling (\mathcal{D}_1)

In the first part of \mathcal{D}_1 , we examine how model performance varies with the degrees of inner and outer polynomials. The result is shown in Figure 1. We use greedy search for the

Experiment	Window	Emb. Dim	Layers	Heads
$\overline{\mathcal{D}_1}$ (Degree)	256	512	6	8
\mathcal{D}_1 (Variable)	850	512	6	8
$\overline{\mathcal{D}_2}$ (Arch.)	850	{256, 512, 768}	{4, 6}	8
\mathcal{D}_2 (Attn.)	850	512	6	{4, 8, 16
$\overline{\mathcal{D}_3}$	256	512	4	8

229

230

231

233

234

235

236

238

239

240

241

242

243

244

245

246

247

248

249

250

Table 2. Common settings: GELU activation, learned positional embeddings, multi-head attention with causal masking, MLP hidden dim = $4 \times$ embedding dim.



Figure 1. Performance across different d_{inner} , d_{outer}

inference. Regardless of the degrees of the polynomials, our model achieves a remarkable single-output accuracy. Notably, when using beam search with a width of 10, the model's accuracy reaches 100% for these configurations.

Our analysis reveals a pattern: performance remains invari-251 ant to increases in the outer polynomial's degree, while 252 253 decreasing when the inner polynomial's degree increases. This demonstrates that the transformer's decomposition ca-254 255 pability is primarily limited by the complexity of the inner 256 polynomial rather than that of the outer polynomial.

257 In the second part of \mathcal{D}_1 , we investigate how the perfor-258 mance scales with v_{inner} and v_{outer} , the number of variables 259 in the inner and outer polynomials. Figures 2 and 3 present these results. 261

Given the challenging nature of multivariate polynomial decomposition, we evaluate the model's performance using 263 beam search with a width of 30, considering a prediction 264 265 correct if at least one of the 30 candidate outputs is correct 266 decomposition.

267 Our results reveal two trends: performance decreases dra-268 matically as v_{outer} increases, yet counter-intuitively im-269 proves as v_{inner} increases. This observation aligns with 270 the following heuristic understanding: higher v_{outer} creates 271 an information bottleneck, requiring the model to simulta-272 neously resolve multiple interdependent inner functions. In 273 contrast, higher v_{inner} provides more dimensions of input 274



Figure 2. Performance across different v_{outer}

Figure 3. Performance across different v_{inner}



Figure 4. Accuracies on different number of layer and dimension.

variation with additional structural indicators that can guide the decomposition process.

4.2. Architecture Scaling (\mathcal{D}_2)

In \mathcal{D}_2 , we examine how model performance varies with architectural parameters: embedding dimension, number of layers, and number of attention heads. When varying the number of heads, we maintain a constant total embedding dimension, meaning that models with more heads have smaller per-head embedding dimensions. We use the dataset described in Section 3.2 and evaluate using beam search with a width of 30.

Figure 4 reveals the scaling behavior (Kaplan et al., 2020) of transformer architectures on polynomial decomposition. As model capacity increases through higher embedding dimensions and additional layers, performance consistently improves.

Notably, our results demonstrate the presence of a datadependent scaling threshold. With limited training data (1M examples), larger models initially underperform their simpler counterparts, particularly evident in the 6-layer configurations with higher embedding dimensions. However, this pattern reverses completely with additional training data, confirming that larger models possess superior capacity for mathematical pattern recognition when provided with



Figure 5. Performance recovery when adapting to a new coefficient distribution

⁵⁸ sufficient examples to leverage their parametric advantage.

289 In \mathcal{D}_2 , we also examine model performance with different 290 numbers of attention heads. Our experiments reveal that 291 increasing the number of attention heads while maintaining 292 constant total embedding dimension leads to progressively 293 deteriorating performance on polynomial decomposition 294 tasks. Models with 4 heads achieved 32.0% accuracy, while 295 those with 8 and 16 heads reached only 28.0% and 25.0% ac-296 curacy, respectively. This suggests that for our specific task 297 of mathematical pattern recognition, fewer, more expressive attention heads with larger per-head dimensions provide 299 better performance than numerous specialized heads with 300 smaller dimensions. 301

4.3. Distribution Adaptation (\mathcal{D}_3)

302

303

304 We evaluate $\mathcal{G}(n)$ as defined in Eq. 5, which measures how 305 quickly models adapt to new coefficient distributions as 306 a function of adaptation sample size n. For this experi-307 ment, we train a model with 4 layers and 512 embedding 308 dimension on the dataset described in Section 3.2. The 309 initial training used 1M examples with outer polynomial 310 coefficient range C_1 , followed by fine-tuning on n examples 311 with coefficient range C_2 for a single epoch. We report the 312 variance in accuracy based on three independent trials. 313

Models trained exclusively on the first dataset achieve only 314 5.67% accuracy on the new distribution, despite reaching 315 nearly 100% accuracy on the original distribution. Figure 5 316 illustrates how performance recovers during adaptation. No-317 tably, despite using only $\approx 2\%$ of the original training data 318 size, the model rapidly recovers its accuracy from single 319 digits to over 90%. This rapid adaptation indicates suc-320 cessful transfer learning, suggesting that the model develops a general mathematical understanding of polynomial 322 substructures rather than memorizing specific numerical 323 324 relationships.

We further investigate whether alternative data representations could enhance this adaptation capability. We propose "split" representation of polynomials, where we randomly select terms from the expanded form and split their coeffi-



Figure 6. Beam scaling withFigure 7. Beam scaling withvarying v_{outer} ($v_{inner} = 3$)varying v_{inner} ($v_{outer} = 3$)

cients. For example:

$$f_{\text{non-split}}(a) = -63 + 23a - 71a^2 - 11a^3 - 14a^4 - 12a^5 - 2a^6$$
$$f_{\text{split}}(a) = -63 + 23a - 4a^2 - 67a^2 - 8a^3 - 3a^3$$
$$-7a^4 - 7a^4 - 12a^5 - a^6 - a^6$$
(6)

In Figure 5, the red line demonstrates $\mathcal{G}(n)$ of the model trained on data with both normal and split representation. Models trained on this mixed data including split representation demonstrate significantly faster adaptation, requiring only 70% of the additional training examples to reach equivalent performance on the new distribution.

This enhanced generalization likely stems from the model being forced to recognize mathematically equivalent but differently represented polynomials, compelling it to develop a deeper understanding of polynomial structure rather than memorizing specific patterns.

4.4. Search Strategy Analysis (\mathcal{D}_4)

We evaluate how search strategies impact model performance on polynomial decomposition tasks, with a particular focus on beam search efficiency. Figure 6 and 7 illustrate the accuracy achieved across different beam widths for polynomials with varying numbers of variables.

Our results reveal an unusually dramatic impact of beam search for polynomial decomposition compared to typical NLP tasks. For two-variable polynomials, accuracy improves from 11% with greedy search to 69% with a beam width of 30—a remarkable 6.3× improvement. This stands in stark contrast to standard neural machine translation applications, where beam search typically yields BLEU score improvements of only 2-4 points (Huang et al., 2018; Ranzato et al., 2016). Even more telling, most NMT systems show diminishing returns with beam widths beyond 5-10 (Freitag & Al-Onaizan, 2017).

4.5. BGRPO Results

We evaluated BGRPO across models of varying sizes from our architecture scaling experiments(D_2), implementing versions both with and without rank signal. Fig 8 illustrates



Figure 8. Accuracies on experiments with different dimension. Each experiment we have finetuned model with 2M data and models trained with BGRPO with and without rank signal on top of that.

these results.

355

357

358 359

360

BGRPO consistently improved accuracy across all beam widths regardless of model size. Without rank signal,
BGRPO gives average accuracy increases of 34.0%, 17.8%,
and 12.4% for 6-layer models with dimension 256, 512, and
768 respectively. Including rank signal in BGRPO produces even more improvements, with average accuracy increases of 46.6%, 28.4%, and 30.2%.

These improvements translate to significant computational 369 efficiency gains. For instance, the dimension-256 model ini-370 tially achieved 26.1% accuracy with beam width 30. After 371 applying BGRPO with rank signal, comparable accuracy (26.0%) was achieved with just beam width 16. This effec-373 tively halves the required beam width for equivalent per-374 375 formance. Since beam search computation scales quadratically with beam width, this improvement reduces beam 376 search computation by approximately 75% while maintain-377 ing equivalent performance. 378

On average, BGRPO without rank signal reduced the required beam width by 31.3%, 14.9%, and 11.4% for 6-layer models with dimension 256, 512, and 768 respectively.
When incorporating rank signal, BGRPO reduced required beam width even further, by 38.9%, 22.0%, and 26.5%.

4.6. Simplification Comparison with Mathematica

While polynomial simplification and polynomial decomposition represent two distinct mathematical objectives, simplification frequently arises as a consequence of decomposition, since decomposed forms generally exhibit reduced algebraic complexity compared to the original expression. In this subsection, we briefly explore the capabilities of our models for this related problem, and benchmark against the most powerful symbolic computation engine Mathematica. Despite our lightweight parameter budgets and the absence of any explicit simplification objective in our training, the models were able to reduce the leaf count (Wolfram Research, Inc., 1996) of complex expressions, with performance on par with — and in two of five complexity regimes surpassing —Mathematica's state-of-the-art FullSimplify function (see Table 3, competitive performances are bolded).

Table 3. Average leaf count comparison (Beam width = 30)

Problem Complexity		Leaf		
v_O	v_S	Transformer	Mathematica	Δ
2	3	27.28	30.03	-2.75
3	3	22.85	22.12	0.73
4	3	22.52	20.00	2.52
3	2	17.27	17.10	0.17
3	4	26.04	27.56	-1.52

These findings highlight that transformers' inherent ability to uncover latent patterns rivals that of the most advanced symbolic computation methods.

5. Conclusion

Our investigation into transformers for polynomial decomposition uncovers key insights into how neural networks can infer hidden algebraic structures.

We find that model performance depends asymmetrically on polynomial complexity parameters (\mathcal{D}_1): inner polynomial degree plays a dominant role, while outer polynomial complexity has limited impact. Counterintuitively, increasing the number of inner variables improves accuracy by imposing structural constraints, whereas more outer variables create information bottlenecks.

From an architectural viewpoint (D_2), we confirm that performance scales with model size. We observe that fewer but more expressive attention heads are especially effective for this task. In terms of distribution adaptation (D_3), models transfer rapidly to new coefficient distributions, requiring as little as 2% of the original training data, indicating that they internalize generalizable principles rather than rely on memorization. Moreover, we can enhance this generalization capability through strategic dataset design.

Beam search analysis (\mathcal{D}_4) yields up to 6.3× improve-

ment over greedy decoding due to the sparse, precise nature of mathematical solutions. Models finetuned with
our rank-aware BGRPO reinforcement learning method
achieve equivalent accuracy with up to 50% smaller beam
widths, cutting inference computation by approximately
75%. Lastly, our model demonstrates competitive performance in polynomial simplification compared with symbolic computation tools in Mathematica.

Our work provides, for the first time, a systematic analysis of transformer capabilities for polynomial decomposition through carefully controlled experiments across four 396 dimensions. Our methodologies can serve as a road map 397 for exploring neural models in other domains that require 398 non-local latent pattern discovery, such as functional decom-399 position problems ranging from systems engineering and 400 mechanical design to digital logic design. While we devel-401 oped BGRPO specifically for enhancing beam search in the 402 polynomial decomposition problem, similar techniques may 403 prove useful in other domains with sparse solution spaces 404 where models can identify correct structures but struggle 405 with specific details. 406

Limitations Our generalizability investigation was constrained to univariate polynomials with relatively narrow coefficient ranges and limited maximum degrees. Computational constraints restricted our architecture scaling experiments to relatively small models (maximum 6 layers, 768-dimensional embeddings); however, the consistent performance improvements without accuracy saturation suggest that further scaling would yield additional gains. Finally, our method's reliance on wide beam search creates computational overhead during inference despite BGRPO's improvements.

References

407

408

409

410

411

412

413

414

415

416

417

418

419

420

421

422

423

424

425

426

427

428

429

430

431

432

433

- Adamski, M. A., Karatkevich, A., Wegrzyn, M., Rawski, M., Łuba, T., Jachna, Z., and Tomaszewicz, P. The influence of functional decomposition on modern digital design process. *Design of Embedded Control Systems*, pp. 193– 204, 2005.
- Dang, T. and Testylier, R. Reachability analysis for polynomial dynamical systems using the bernstein expansion. *Reliab. Comput.*, 17(2):128–152, 2012.
- DeepSeek-AI. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning, 2025. URL https://arxiv.org/abs/2501.12948.
- Demirtas, S., Su, G., and Oppenheim, A. V. Sensitivity of polynomial composition and decomposition for signal processing applications. In 2012 Conference Record of the Forty Sixth Asilomar Conference on Signals, Systems and Computers (ASILOMAR), pp. 391–395. IEEE, 2012.

- Dickerson, M. T. Polynomial decomposition algorithms for multivariate polynomials. Technical report, Cornell University, 1987.
- Dickerson, M. T. General polynomial decomposition and the s-1-decomposition are np-hard. *International Journal of Foundations of Computer Science*, 4(02):147–156, 1993.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv* preprint arXiv:2010.11929, 2020.
- Drori, I., Zhang, S., Shuttleworth, R., Tang, L., Lu, A., Ke, E., Liu, K., Chen, L., Tran, S., Cheng, N., et al. A neural network solves, explains, and generates university math problems by program synthesis and few-shot learning at human level. *Proceedings of the National Academy of Sciences*, 119(32):e2123433119, 2022.
- Elias, A. J. and Wen, J. T. Ik-geo: Unified robot inverse kinematics using subproblem decomposition. *Mechanism and Machine Theory*, 209:105971, 2025.
- Faugère, J.-C. and Perret, L. An efficient algorithm for decomposing multivariate polynomials and its applications to cryptography. *Journal of Symbolic Computation*, 44 (12):1676–1689, 2009a.
- Faugère, J.-C. and Perret, L. High order derivatives and decomposition of multivariate polynomials. In *Proceedings* of the 2009 international symposium on Symbolic and algebraic computation, pp. 207–214, 2009b.
- Fawzi, A., Kozhasov, K., Goldblum, M., Behrmann, J., Zhang, C., Fuchs, F., Huang, P.-S., Li, L., and Kohli, P. Discovering faster matrix multiplication algorithms with reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2022.
- Freitag, M. and Al-Onaizan, Y. Beam search strategies for neural machine translation. In *Proceedings of the First Workshop on Neural Machine Translation*. Association for Computational Linguistics, 2017. doi: 10.18653/v1/w17-3207. URL http://dx.doi.org/ 10.18653/v1/W17-3207.
- Gathen, J. v. z., Gutierrez, J., and Rubio, R. Multivariate polynomial decomposition. *Applicable Algebra in Engineering, Communication and Computing*, 14(1):11–31, 2003.
- Hernandez, I., Watson, B. C., Weissburg, M. J., and Bras, B. Using functional decomposition to bridge the design gap between desired emergent multi-agent-system resilience and individual agent design. *Systems Engineering*, 27(5): 911–930, 2024.

- Huang, L., Zhao, K., and Ma, M. When to finish? optimal beam search for neural text generation (modulo
 beam size), 2018. URL https://arxiv.org/abs/
 1809.00069.
- Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., Tunyasuvunakool, K., Bates, R., Žídek, A., Potapenko, A., et al. Highly accurate protein structure prediction with alphafold. *nature*, 596(7873):583–589, 2021.
- Kaplan, J., McCandlish, S., Henighan, T., and Brown, T. B.
 Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.
- 454 Karpathy, A. mingpt. https://github.com/455 karpathy/minGPT, 2020.

456

457

458

463

464

465

466

467

- Lample, G. and Charton, F. Deep learning for symbolic mathematics. *arXiv preprint arXiv:2006.02974*, 2020.
- Lin, H.-P., Jiang, J.-H. R., and Lee, R.-R. To sat or not to
 sat: Ashenhurst decomposition in a large scale. In 2008 *IEEE/ACM International Conference on Computer-Aided Design*, pp. 32–37. IEEE, 2008.
 - Liu, Z., Chen, C., Li, W., Qi, P., Pang, T., Du, C., Lee, W. S., and Lin, M. Understanding r1-zero-like training: A critical perspective, 2025. URL https://arxiv. org/abs/2503.20783.
- Manocha, D. and Canny, J. F. Real time inverse kinematics
 for general 6r manipulators. In *ICRA*, pp. 383–389, 1992.
- Mori, M., Cheng, C., Taylor, B. R., Okano, H., and Hwa,
 T. Functional decomposition of metabolism allows a
 system-level quantification of fluxes and protein allocation towards specific metabolic functions. *Nature Com- munications*, 14(1):4161, 2023.
- 476 Olsson, C., Elhage, N., Nanda, N., Joseph, N., DasSarma, 477 N., Henighan, T., Mann, B., Askell, A., Bai, Y., Chen, 478 A., Conerly, T., Drain, D., Ganguli, D., Hatfield-Dodds, 479 Z., Hernandez, D., Johnston, S., Jones, A., Kernion, J., 480 Lovitt, L., Ndousse, K., Amodei, D., Brown, T., Clark, J., 481 Kaplan, J., McCandlish, S., and Olah, C. In-context learn-482 ing and induction heads. Transformer Circuits Thread, 483 2022. https://transformer-circuits.pub/2022/in-context-484 learning-and-induction-heads/index.html. 485
- Patarin, J. and Goubin, L. Asymmetric cryptography with
 s-boxes is it easier than expected to design efficient asymmetric cryptosystems? In *International Conference on Information and Communications Security*, pp. 369–380.
 Springer, 1997.
- Polu, S. and Sutskever, I. Generative language modeling for automated theorem proving. *arXiv preprint arXiv:2009.03393*, 2020.

- Ranzato, M., Chopra, S., Auli, M., and Zaremba, W. Sequence level training with recurrent neural networks, 2016. URL https://arxiv.org/abs/1511. 06732.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms, 2017. URL https://arxiv.org/abs/ 1707.06347.
- Shao, Z., Wang, P., Zhu, Q., Xu, R., Song, J., Bi, X., Zhang, H., Zhang, M., Li, Y. K., Wu, Y., and Guo, D. Deepseekmath: Pushing the limits of mathematical reasoning in open language models, 2024. URL https://arxiv.org/abs/2402.03300.
- She, J., Belanger, E., and Bartels, C. Evaluating the effectiveness of functional decomposition in early-stage design: development and application of problem space exploration metrics. *Research in Engineering Design*, 35 (3):311–327, 2024.
- Tempero, E., Denny, P., Finnie-Ansley, J., Luxton-Reilly, A., Kirk, D., Leinonen, J., Shakil, A., Sheehan, R., Tizard, J., Tu, Y.-C., et al. On the comprehensibility of functional decomposition: An empirical study. In *Proceedings of the* 32nd IEEE/ACM International Conference on Program Comprehension, pp. 214–224, 2024.
- Trinh, T. H., Wu, Y., Le, Q. V., He, H., and Luong, T. Solving olympiad geometry without human demonstrations. *Nature*, 625(7995):476–482, 2024.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. *Advances in neural information* processing systems, 30, 2017.
- von Werra, L., Belkada, Y., Tunstall, L., Beeching, E., Thrush, T., Lambert, N., Huang, S., Rasul, K., and Gallouédec, Q. Trl: Transformer reinforcement learning. https://github.com/huggingface/trl, 2020.
- Von Zur Gathen, J. Functional decomposition of polynomials: the tame case. *Journal of Symbolic Computation*, 9 (3):281–299, 1990a.
- Von Zur Gathen, J. Functional decomposition of polynomials: the wild case. *Journal of Symbolic Computation*, 10 (5):437–452, 1990b.
- Wang, K., Variengien, A., Conmy, A., Shlegeris, B., and Steinhardt, J. Interpretability in the wild: a circuit for indirect object identification in gpt-2 small, 2022. URL https://arxiv.org/abs/2211.00593.

- 495 Wolfram Research, Inc. ComplexityFunction
 496 Wolfram Language Documentation. https:
 497 //reference.wolfram.com/language/ref/
 498 ComplexityFunction.html, 1996. Accessed
 499 12 May 2025.
- 501 Zhao, S., Feng, R., and Gao, X.-S. On functional decomposition of multivariate polynomials with differentiation
 503 and homogenization. *Journal of Systems Science and*504 *Complexity*, 25(2):329–347, 2012.

A. Backward Synthetic Data Generation Algorithm

Our backward synthetic data generation in subsection 2.1 can be described as follows.

 513
 Algorithm 1 Backward Generation of Synthetic Training

 514
 Data

- **Require:** Coefficient range C_{inner} , C_{outer} ; maximal degrees d_{inner} , d_{outer} ; variable counts v_{inner} , v_{outer} ; term limits t_{inner} , t_{outer} .
- 1: Generate outer polynomial g with v_{outer} variables, coefficients $\in C_{outer}$, degree $= d_{outer}$, and no more than t_{outer} monomial terms.
- 2: Generate v_{outer} inner polynomials $h_1, \ldots, h_{v_{\text{outer}}}$, where each h_i has v_{inner} variables, coefficients $\in C_{\text{inner}}$, degree = d_{inner} , and no more than t_{inner} monomial terms.
 - 3: $f \leftarrow g(h_1, \ldots, h_{v_{\text{outer}}})$, i.e. substitute $h_1, \ldots, h_{v_{\text{outer}}}$ into g, expand and collect the monomial terms.
 - 4: **return** $(f, g, h_1, ..., h_{v_{outer}})$

B. Tokenization

500

505 506

507 508

509

510

511 512

515

516

517

518

519

520

521

522

523

524

525

526

527

528

529

530

531

We encode polynomials using prefix notation, with separate 532 tokens for operators, digits, and variables. Each number in-533 cludes its sign, so we only use addition, multiplication, and 534 power operators. Subtraction is represented as addition with 535 a negative sign. Each input sequence consists of the tok-536 enized expanded polynomial f followed by a question mark 537 token '?'. The target output format depends on the number 538 of outer variables: for $v_{outer} = 1$, the target output is simply 539 the tokenized inner polynomial h; for $v_{outer} > 1$, the target 540 output begins with the tokenized outer polynomial g fol-541 lowed by each tokenized inner polynomial $h_1, \ldots, h_{v_{outer}}$, 542 with all polynomials separated by a delimiter token '&'. 543

Below is an example of a tokenized training input 'x' and target output 'y':



Figure 9. Top-3 probability for each token position in the answer sequence where

Answer: + * N 5 ^ b1 P 3 + * N 4 * b0 ^ b2 P 2 * N 5 ^ b2 P 3 & + * P 2 * ^ a0 P 2 a2 * N 2 * a0 * a1 a2 & + * N 5 ^ a0 P 3 + * P 4 * ^ a0 P 2 a2 * N 5 * a0 * a1 a2 & + * P 4 * a0 * a1 a2 * P 2 * a1 ^ a2 P 2

 $\begin{array}{l} \textbf{Question:} + * P \ 6 \ 2 \ 5 \ \ a0 \ P \ 9 \ + * \ N \ 1 \ 5 \ 0 \ 0 \ * \ \ a0 \ P \ 8 \\ a2 \ + * P \ 1 \ 8 \ 7 \ 5 \ * \ \ a0 \ P \ 7 \ \ a1 \ a2 \ + * P \ 1 \ 2 \ 0 \ 0 \ * \ \ a0 \ P \ 7 \ \ a2 \ P \ 2 \\ + * \ N \ 3 \ 0 \ 0 \ 0 \ * \ \ a0 \ P \ 6 \ \ a1 \ \ a2 \ P \ 2 \ + \ P \ 1 \ 8 \ 7 \ 5 \ * \ \ a0 \ P \ 5 \ \ * \ a0 \ P \ 5 \ \ a0 \ \ a) \ a0 \ \ a0 \ \ a) \ a0 \ \ a0 \$

This example shows a training pair where the outer

polynomial is $90a - 319a^2 - 36a^3 - a^4$ and the target inner polynomial is $-5 + 18a + a^2$. The \Box symbol represents a padding token which is excluded from the log-likelihood loss calculation.

C. Example Ouput Logits and Effectiveness of the Beam Search

Figure 9 shows example top-3 probabilities for each token position in the answer sequence at temperature 1, using the layer-6, embedding dimension 512 model from our D_2 experiments. Correct answers are highlighted in red. The visualization clearly illustrates that the model's primary source of confusion occurs in sign decisions, while it confidently predicts most of the other token types.

Table 4 quantifies this observation by showing the probability and accuracy statistics for different token types across our model architectures from D_2 . These statistics were computed using a test set of 1000 polynomial decomposition problems at temperature 1.

557 *Table 4.* Token Type Analysis Across Different Model Architec-

556

564

565

566 567

585

586

Token Type	Metric	4 Layers			6 Layers		
		256 dim	512 dim	768 dim	256 dim	512 dim	768 din
Sign	Probability Accuracy	$\begin{array}{c} 0.489 \pm 0.001 \\ 0.519 \pm 0.006 \end{array}$	$\begin{array}{c} 0.489 \pm 0.001 \\ 0.531 \pm 0.006 \end{array}$	$\begin{array}{c} 0.493 \pm 0.001 \\ 0.530 \pm 0.006 \end{array}$	$\begin{array}{c} 0.491 \pm 0.001 \\ 0.522 \pm 0.006 \end{array}$	$\begin{array}{c} 0.490 \pm 0.001 \\ 0.523 \pm 0.006 \end{array}$	$0.490 \pm 0.$ $0.521 \pm 0.$
Operator	Probability Accuracy	$\begin{array}{c} 0.920 \pm 0.002 \\ 0.937 \pm 0.002 \end{array}$	$\begin{array}{c} 0.915 \pm 0.002 \\ 0.934 \pm 0.002 \end{array}$	$\begin{array}{c} 0.919 \pm 0.002 \\ 0.935 \pm 0.002 \end{array}$	$\begin{array}{c} 0.927 \pm 0.002 \\ 0.943 \pm 0.002 \end{array}$	$\begin{array}{c} 0.925 \pm 0.002 \\ 0.941 \pm 0.002 \end{array}$	$0.925 \pm 0.0000000000000000000000000000000000$
Number	Probability Accuracy	$\begin{array}{c} 0.880 \pm 0.002 \\ 0.901 \pm 0.002 \end{array}$	$\begin{array}{c} 0.870 \pm 0.002 \\ 0.893 \pm 0.003 \end{array}$	$\begin{array}{c} 0.878 \pm 0.002 \\ 0.897 \pm 0.002 \end{array}$	$\begin{array}{c} 0.890 \pm 0.002 \\ 0.911 \pm 0.002 \end{array}$	$\begin{array}{c} 0.885 \pm 0.002 \\ 0.905 \pm 0.002 \end{array}$	$0.884 \pm 0.$ $0.903 \pm 0.$

Table 4. Note: Values shown as mean \pm standard error of the mean. The sign token probabilities are near-random, while operators and numbers show high confidence and accuracy.

As discussed in Section 2.2, our models achieve approximately 90% accuracy when predicting non-sign tokens, but
exhibit near-random performance when choosing between
positive and negative signs. This specific error pattern makes
beam search particularly effective for our task.

573 The effectiveness of beam search stems from its ability to 574 explore multiple sign configurations while preserving the 575 high-confidence structural tokens. In probability terms, se-576 lecting a token with 0.1 probability instead of one with 0.9 577 probability is equivalent to making approximately 11 con-578 secutive choices of a 0.45 probability token over a 0.55 prob-579 ability token. Since our polynomial expressions typically 580 contain fewer than 10 sign decisions, beam search with a 581 width of approximately 30 can efficiently cover most viable 582 sign permutations while maintaining the correct monomial 583 structure identified with high confidence. 584

D. Attention Score Analysis: Monomial Heads

Attention mechanism analysis has provided valuable insights into transformer model behaviors, with studies identifying specialized attention heads that serve specific functions. For example, (Olsson et al., 2022) identified "Induction Heads" that play a crucial role in in-context learning,
while (Wang et al., 2022) provided a comprehensive understanding of indirect object identification in GPT-2 Small.

595 In our analysis of attention patterns in polynomial decompo-596 sition models, we identified specialized attention heads that 597 recognize the structure of polynomials, particularly focus-598 ing on monomial identification. We call these "Monomial 599 Heads," and they appear consistently across all model sizes 600 in our architecture scaling experiments (D_2).

Monomial Heads manifest in two distinct patterns in our
 models. First, in layer 0, several attention heads consistently
 attend to tokens 1-5 positions behind the current position,

as shown in the leftmost plot of Figure 10. Second, in layer 1, we observe specialized behavior where certain heads focus attention on specific tokens within each monomial of the input polynomial (middle plot), while others specifically attend to delimiter tokens in the decomposition output (rightmost plot).

We hypothesize that this represents a two-stage process: in the first layer, the model identifies key tokens that serve as indicators for each monomial by examining local context (1-5 tokens behind). In the second layer, tokens within each monomial attend to these indicator tokens to establish their monomial membership. While this pattern is most clear in the encoding of the input polynomial, the decomposition output shows evidence of boundary recognition, particularly at the transitions between inner functions marked by delimiter tokens.



Figure 10. Attention score visualization of selected attention heads from our 6-layer transformer model with embedding dimension 768. The visualization shows attention patterns for a tokenized polynomial sequence and its decomposition.

 $\begin{array}{l} \textbf{Model's decomposition output:} & + * N \ 4 \ \wedge \ b0 \ P \ 3 \ + * \\ b0 \ \wedge \ b2 \ P \ 2 \ * N \ 1 \ \wedge \ b2 \ P \ 3 \ \& + * N \ 4 \ \wedge \ a0 \ P \ 3 \ * \ \wedge \ a0 \ P \ 2 \ a1 \ \& \\ & + * N \ 3 \ \wedge \ a1 \ P \ 3 \ + * N \ 2 \ * \ a1 \ \wedge \ a2 \ P \ 3 \ \& \ * N \ 4 \ \wedge \ a1 \ P \ 3 \end{array}$

The visualization reveals how different attention heads focus on specific structural elements when decomposing polynomials.