A Appendix

A.1 Code availability

The source code underlying this project is available at the GitHub repository `https://github.com/Thorben010/llm_synthesis`.

A.2 Additional results

Table 3: **Model comparison.** Comparison of embedding methods on different data regimes for sintering and calcination temperatures. We report mean across five runs with standard deviation in parentheses. ✓ indicates training on the respective data source.

	Synth. Data	Literature Data	Sintering temperature			Calcination temperature			Rel. MAE Imp.
Model			**MAE ↓**	**RMSE ↓**	**R^2 ↑**	**MAE ↓**	**RMSE ↓**	**R^2 ↑**	↑
SyntMTE	✓	✓	**135.00** (0.84)	**181.30** (0.71)	**0.545** (0.004)	**153.72** (0.54)	**199.71** (0.48)	**0.436** (0.003)	4.08%
SyntMTE		✓	141.00 (2.13)	189.27 (2.46)	0.504 (0.013)	160.00 (2.97)	206.56 (2.53)	0.395 (0.015)	0.00%
SyntMTE	✓		149.43 (3.53)	197.78 (1.47)	0.428 (0.014)	169.63 (2.46)	214.64 (1.87)	0.358 (0.019)	-6.00%
CrabNet	✓	✓	148.03 (1.00)	196.88 (1.60)	0.464 (0.009)	159.67 (0.89)	206.50 (1.13)	0.397 (0.007)	3.77%
CrabNet		✓	152.87 (6.59)	205.37 (8.20)	0.416 (0.047)	166.88 (3.13)	215.80 (0.46)	0.340 (0.025)	0.00%
CrabNet	✓		160.41 (2.12)	199.83 (0.69)	0.402 (0.016)	172.54 (4.07)	216.66 (1.47)	0.329 (0.036)	-4.13%
Composition + NN	✓	✓	149.75 (0.87)	191.68 (0.78)	0.492 (0.004)	162.82 (0.41)	208.44 (0.50)	0.385 (0.003)	0.96%
Composition + NN		✓	150.23 (2.43)	193.82 (3.21)	0.480 (0.017)	165.38 (3.41)	211.54 (3.74)	0.366 (0.022)	0.00%
Composition + NN	✓		170.58 (2.74)	203.32 (0.06)	0.339 (0.016)	176.87 (3.56)	219.88 (0.00)	0.315 (0.017)	-10.09%
Composition + XGBoost	✓	✓	163.62 (0.59)	210.47 (0.73)	0.387 (0.004)	179.54 (0.84)	228.00 (1.07)	0.263 (0.007)	-13.23%
Composition + XGBoost		✓	141.12 (0.90)	189.14 (1.45)	0.505 (0.008)	161.96 (0.87)	206.65 (1.05)	0.395 (0.006)	0.00%
Composition + XGBoost	✓		196.03 (1.90)	242.25 (1.83)	0.188 (0.012)	225.74 (3.49)	276.76 (3.48)	-0.086 (0.027)	-39.16%

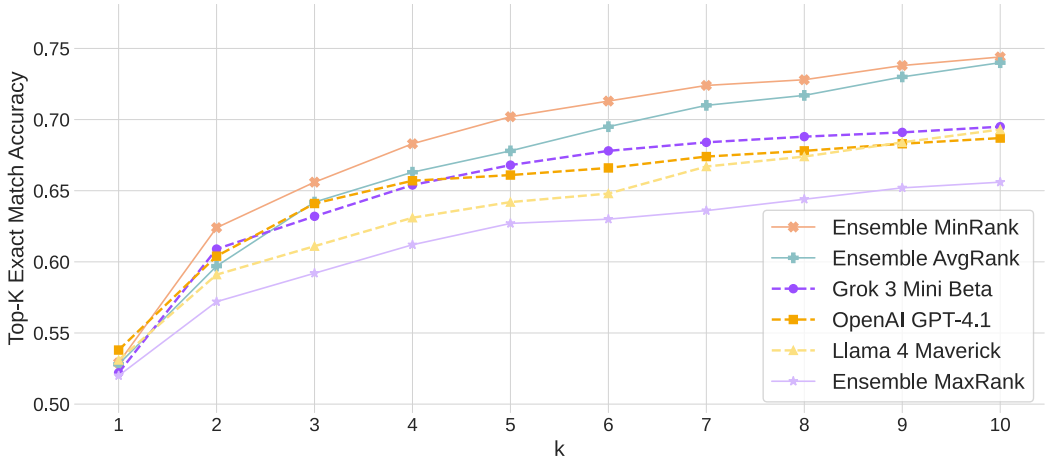

Figure 5: **LM ensemble comparison.** Top-k exact-match accuracy for three individual language models: Grok 3 Mini, GPT-4.1, and Llama 4 Maverick; and their joint ensemble with predictions combined using minimum-rank, average-rank, and maximum-rank voting. The minimum-rank ensemble achieves the best recall beyond Top–1.

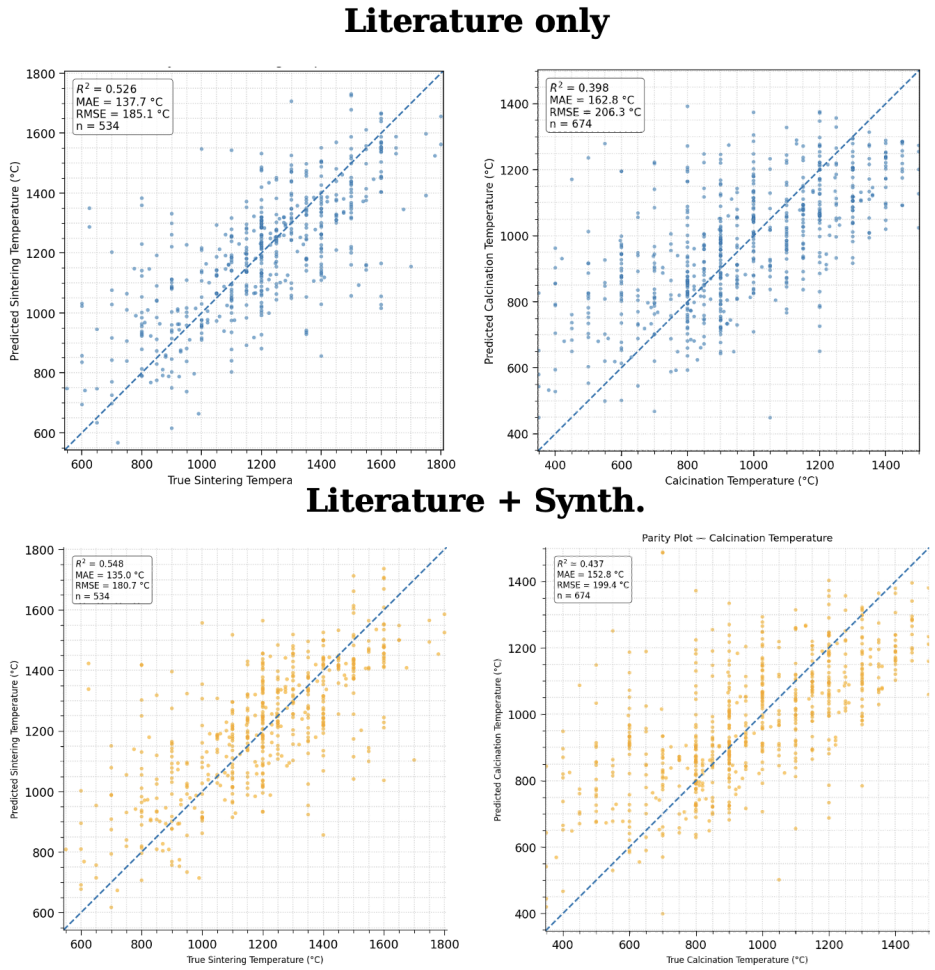

Figure 6: Parity plots for the regression task shown in Table 3. The literature only model has been trained on literature only, while the model Literature + Synth. is sequentially trained on both datasets. Notably R^2 scores are increased in the second setting.

A.3 Dataset

We base our study on the corpus introduced by Kononova *et al.* [16], containing 33,343 solid-state synthesis records extracted from the literature with paragraph- and phrase-level NLP models that tag targets, precursors, and by-products. Following quality control, removal of entries with unclear stoichiometry and enforcement of element balance between targets and their precursor sets, the working set comprises 18,804 reactions, of which 9,255 are unique. For evaluation, we use the year-stratified, Distinct Reactions, and Novel Material Systems (NMS) splits described by Prein *et al.* [26].

A.4 LM overview

In our study, we systematically evaluated seven contemporary LMs, selected to represent a diverse array of architectures, parameter scales and licensing schemes. The main characteristics of these models are summarized in Table 4.

Model	Release	#Params	#Active	Context Window	Open Source	ELO Score	MMLU-Pro
Qwen 2.5 VL-72B	Jan 2025	72B	-	32k tokens	Yes (Qwen License)	1123	71.2
Mistral Small 3.1	Mar 2025	24B	-	128k tokens	Yes (Apache 2.0)	1249	66.8
DeepSeek-V3-0324	Mar 2025	671B	37B	128k tokens	Yes (MIT License)	1369	81.2
Gemini 2.0 Flash	Feb 2025	-	-	1M tokens	No	1352	76.4
GPT-4.1	Apr 2025	-	-	1M tokens	No	1365	-
LLaMA 4 Maverick	Apr 2025	400B	17B	1M tokens	Yes (LLaMA 4 Comm. License)	1266	80.5
Grok 3 Mini Beta	Apr 2025	-	-	131k tokens	No	-	78.9

Table 4: LMs evaluated in this work, sorted by release date. Arena scores from the Chatbot Arena leaderboard as of May 18, 2025

Qwen 2.5 VL-72B. Released in January 2025, Qwen 2.5 VL-72B (Apache 2.0 license) comprises 72 billion parameters and features a context window of 32,000 tokens. The model achieved an ELO rating of 1123 and an MMLU-Pro score of 71.2. [70, 71]

Mistral Small 3.1. This model was introduced in March 2025 under the Apache 2.0 license and contains 24 billion parameters. It offers an extended context window of 128,000 tokens. While detailed training data information remains undisclosed, it attains an ELO score of 1249 and an MMLU-Pro score of 66.8, indicating robust generalization abilities. [72, 71]

DeepSeek-V3-0324. This model, released under the MIT license, is a Mixture-of-Experts (MoE) language model featuring 671 billion total parameters, with 37 billion activated per token. Its MoE layers consist of 1 shared and 256 routed experts, with 8 routed experts actively engaged for each token. It's pre-trained on 14.8 trillion tokens and supports a context window extended to 128,000 tokens. The March 2025 checkpoint of this model achieves the highest ELO (1369) and MMLU-Pro (81.2) score among the evaluated models. [73, 74, 71]

Gemini 2.0 Flash. Released in February 2025 as a proprietary model, Gemini 2.0 Flash features a context window of 1 million tokens. While the specific parameter count and training details are not publicly disclosed, the model's performance is notable, reflected by an ELO rating of 1352 and an MMLU-Pro score of 76.4. [71]

GPT-4.1. This model was introduced in April 2025, also supporting a large context window of 1 million tokens. Despite undisclosed training and parameter specifics, it achieved a remarkably high ELO score of 1365. [75, 71]

LLaMA 4 Maverick. The model, released in April 2025 under the LLaMA 4 Community License, incorporates a Mixture-of-Experts (MoE) architecture with 128 experts. It features a grand total of 400 billion parameters, of which approximately 17 billion are actively engaged during inference. Trained on an extensive dataset of roughly 22 trillion tokens, its 1 million token context window supports advanced in-context learning capabilities. Evaluation metrics include an ELO score of 1266 and an MMLU-Pro score of 80.5. [76, 71]

Grok 3 Mini Beta. This proprietary LM, launched in April 2025, provides a context window of 131,000 tokens. It attained an MMLU-Pro score of 78.9. The ELO score was not disclosed publicly. [77]

A.5 LM multi-Provider Inference via OpenRouter Proxy

All model inference is performed through the OpenRouter API, which federates requests to multiple upstream providers under a single authentication and billing framework. For every model we set the temperature parameter at $\tau = 0.1$, giving token probabilities

$$p_\tau(i) = \frac{\exp(z_i/\tau)}{\sum_j \exp(z_j/\tau)} \tag{1}$$

and retained all other provider defaults. All models were benchmarked via OpenRouter on 1000 held-out targets. These models were selected to represent a diverse cross-section of leading commercially available large language models, with a primary focus on balancing strong performance and cost-effectiveness.

Model	Input ($/1M)	Output ($/1M)
GPT-4.1 (OpenAI)	2.0	8.0
Grok 3 Mini Beta (xAI)	0.3	0.5
Llama 4 Maverick (Meta)	0.17	0.85
DeepSeek Chat v3 (DeepSeek)	0.27	1.1
Mistral Small 3.1 (Mistral)	0.10	0.30
Gemini 2.0 Flash-001 (Google)	0.10	0.40
Qwen 2.5 VL 72B (Alibaba)	0.70	0.70

Table 5: Per-token API costs via OpenRouter (May 23, 2025)

All API requests and responses are recorded in a structured log for full traceability. If a response doesn't satisfy the prescribed output format (missing the required list or dictionary structure), the identical prompt is retried up to two additional times, for a total of three format validation attempts. Should all three attempts yield structurally invalid output, the response is marked as a failure. API-level errors are immediately assigned a value of `None`. They are logged, and retried without counting against the three allowed format attempts. This separation ensures that transient infrastructure issues do not penalize model performance estimates.

A.6 LM prompting

In-context learning refers to a language model's ability to solve a novel task from prompt-side demonstrations without parameter updates. Performance in this setting typically improves with the number of examples (or "shots") provided, at least until the model's context window is exhausted. This phenomenon was first highlighted in GPT-3, where accuracy on benchmarks such as SuperGLUE increased monotonically as the number of demonstrations was varied from 1–32, demonstrating that richer prompts can substitute for task-specific fine-tuning [78]. Here, we adopt this methodology to determine an appropriate number of examples for our models. Our evaluation using Mistral Small on the precursor suggestion task shows that overall performance improves up to 40 in-context examples and plateaus beyond this point. Accordingly, we use 40 examples in our evaluations on all models.

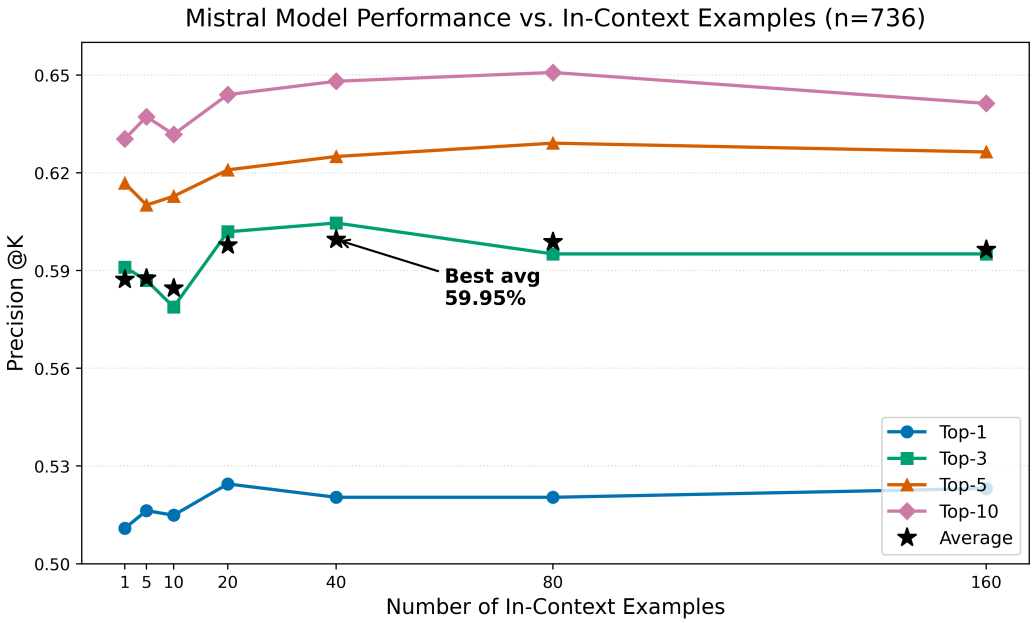

Figure 7: Influence of the number of in-context examples on the performance across 736 samples of the validation set.

A.7 LM evaluation protocols

For the two main tasks at hand: Precursor suggestion and synthesis condition regression, we use the evaluation protocols described below.

A.7.1 LM precursor suggestion task

We evaluate on a 1,000-sample subset of the curated corpus by Kononova *et al.* [16]. For the precursor suggestion task, each LM is prompted to propose feasible precursor sets for a target material. Prompts include 40 few-shot, solid-state examples to fix the output schema. Models return a Python object that we parse into a ranked list of precursor sets (the list index defines the rank). All formulas are then canonicalized to pretty formulas with `pymatgen` [79] prior to scoring.

Our evaluation follows prior work [25, 41]. We iterate over the test set assuming a single ground-truth precursor set per target (the dataset is de-duplicated with respect to target–precursor pairs) and do not merge alternative literature routes for the same target, avoiding bias toward highly reported systems. Exact-match accuracy is reported as the fraction of cases where the ground-truth set appears within the Top-k predictions (Top-1, Top-3, Top-5, *etc.*).

A.8 LM synthesis-condition regression task

For condition prediction, we ask each model to output calcination and sintering temperatures along with the associated dwell times. The test set comprises 1,000 entries from the Kononova corpus [16]. Prompts are augmented with 40 few-shot examples drawn from the validation split (none overlap the test set). Generated responses are post-processed to extract numeric values and scored against ground truth using R^2, mean absolute error (MAE), and root-mean-square error (RMSE). Figure 8 illustrates sintering-temperature performance for Gemini 2.0 Flash, with each point representing one sample.

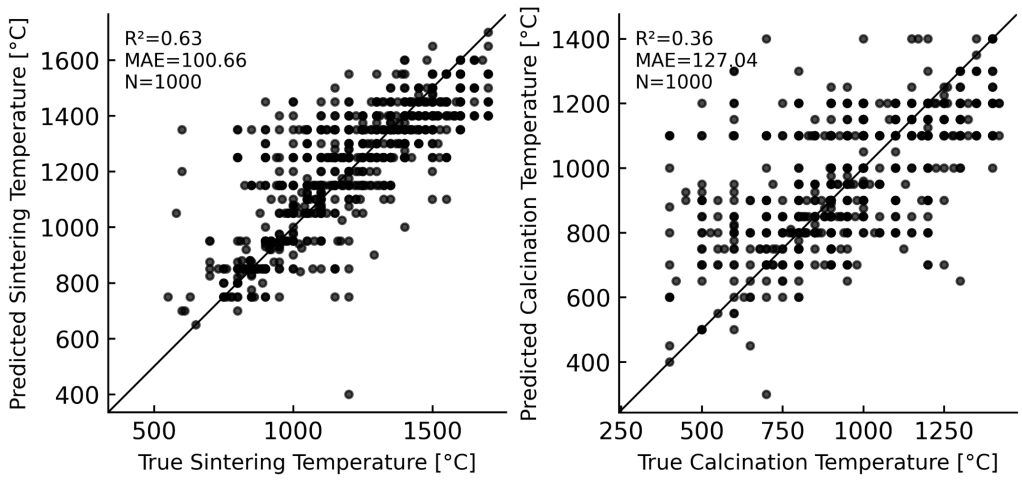

Figure 8: Scatter plot of true versus predicted sintering temperatures for the Gemini 2.0 Flash model on the 40-shot test set with fixed precision length.

A.9 LM additional dataset generation details

In addition to the procedures described in the main text, we retrieved all 48,927 reported syntheses from the Materials Project database [6], filtering exclusively for compounds with documented experimental syntheses. Each chemical formula was parsed into a `Composition` object in `pymatgen` and converted into an atomic-fraction vector over the union of all elements. These high-dimensional vectors were discretized into $M = 1,000$ clusters via MiniBatchKMeans, producing a histogram representation of compositional regions. To select a uniform, maximally diverse subset of $K = 10,000$ formulas, we employed a greedy Shannon-entropy maximization algorithm: at each iteration, a random candidate subset was sampled, the increase in histogram entropy for each candidate was computed, and the candidate yielding the greatest gain was selected. We then prompted GPT-4.1

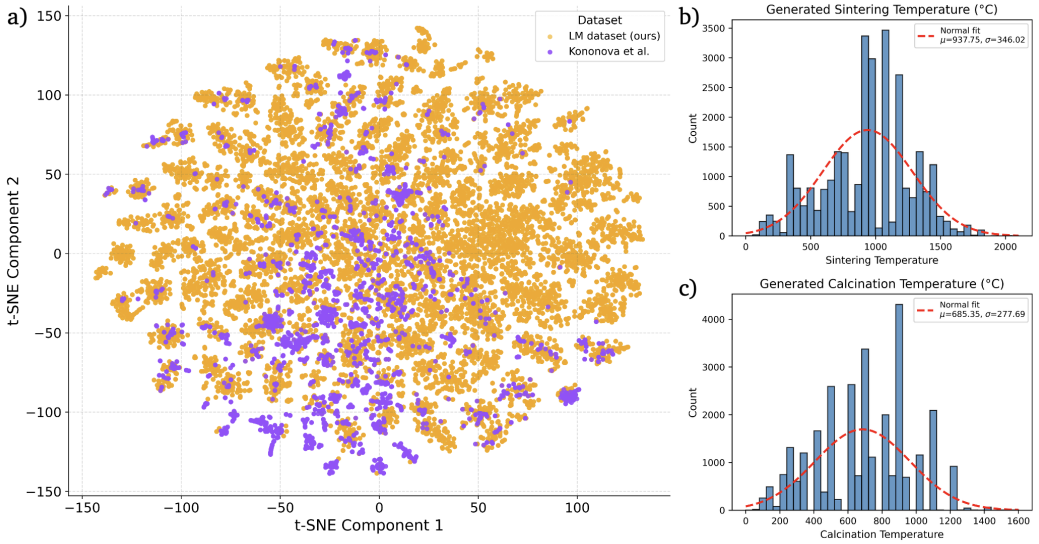

Figure 9: **t-SNE projection of inorganic precursor compositions in embedding space.** a) Compositions from the Kononova dataset [16] are shown in purple, and LM-generated compositions in orange. Compositions are represented by standardized elemental-fraction vectors and projected via t-SNE. The expanded spread of orange points indicates that our generated dataset spans a larger chemical composition space than the baseline dataset. b)–c) Distributions of generated processing parameters.

to assess solid-state synthesizability, flagged compositions were excluded and only those amenable to solid-state methods were retained. For each remaining material, the top three precursor routes, reflecting the model's 64.1 % Top-3 accuracy, were preserved (Table 1). Synthesis-condition parameters were subsequently predicted by GPT-4.1 (Table 2), yielding 29,473 entries. The initial review revealed unphysically low temperatures arising from liquid-phase route suggestions (e.g., $Cu(NH_3)_4(NO_3)_2$ via crystallization rather than solid-state synthesis [80]). To ensure plausibility, we imposed minimum temperature thresholds of 300 °C for calcination and 500 °C for sintering, resulting in the final set of 28,548 solid-state synthesis recipes.

A.10 SyntMTE Model Architecture and Training

SyntMTE is a transformer-based model derived from the MTEncoder framework [41], pretrained on the Alexandria DFT database [7] across 12 materials properties (Table 6). It encodes each reaction by processing the target composition and all precursor materials with shared MTEncoder weights; the resulting embeddings are mean-pooled and concatenated, then passed to a two-layer MLP head for multi-task regression of calcination and sintering temperatures. We fine-tune all weights. Training uses Adam (learning rate 4.39×10^{-5}) with L1 loss, batch size 25, and 200 epochs; the encoder hidden dimension is 512. Experiments run on two NVIDIA RTX A6000 GPUs.

A.11 SyntMTE LLZO case study methodology

We study processing temperatures for LLZO garnet solid electrolytes. Reference LLZO literature is drawn from the corpus compiled by [54]. To ensure strict extrapolation, we exclude from training and validation any record whose target mentions LLZO or a commonly doped variant (Al, Ga, Ta, Nb, W). Calcination and sintering temperatures were extracted using OpenAI's O3 model, followed by manual spot checks. The `SyntMTE` model used for evaluation was fine-tuned sequentially on our synthetic recipes and then on the literature-mined corpus [16], after which it was applied to the mined LLZO dataset. In Fig. 4b, error bars aggregate across distinct literature routes and precursor choices; they reflect across-route variability and do not represent model uncertainty.

A.12 SyntMTE baseline Models

For the regression task we employ several baseline models.

A.12.1 CrabNet

CrabNet is a composition-only materials prediction framework that leverages a transformer encoder to learn contextualized embeddings for each element in a compound [55]. By combining learned element vectors with sinusoidal "fractional embeddings" of stoichiometry, CrabNet's multi-head self-attention layers capture complex inter-element interactions without any hand-crafted descriptors or structural information. This design not only yields state-of-the-art accuracy on benchmarks like MatBench, often outperforming graph-based models such as Roost, [81] but also enables interpretability: attention maps highlight which element pairs drive a given property prediction. In our experiments we use three transformer blocks.

A.12.2 MTEncoder (`SyntMTE`)

Figure 10 depicts the MTEncoder workflow, demonstrating how a material's elemental representation is encoded via a transformer-based model [41]. Each material is broken down into individual element tokens (e.g., Na, Fe, O) alongside a dedicated "Compound" token (*CPD*) that aggregates the element-specific information. These tokens are fed into the transformer encoder, which produces context-rich embeddings for the composition. The embedding associated with the CPD token serves as the learned representation of the material and is passed to an MLP head to predict various properties. Pretraining is conducted using the Alexandria database on 12 tasks (Table 6[7]).

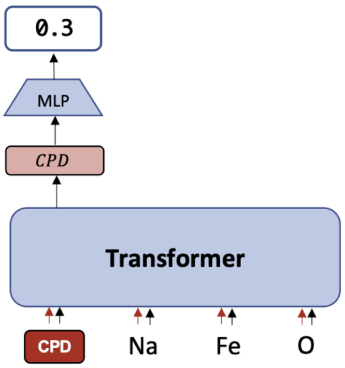

Figure 10: **Overview of the MTEncoder pipeline.** Material compositions are tokenized and processed by a transformer to generate feature embeddings for downstream property prediction.

Pretraining Objectives
Stress
Band Gap (Direct)
Band Gap (Indirect)
Density of States at Fermi Level
Energy Above Hull
Formation Energy
Corrected Total Energy
Phase Separation Energy
Number of Atomic Sites
Total Magnetic Moment
Crystal Space Group
Masked Element Reconstruction (Self-Supervised)

Table 6: **Pretraining objectives for MTEncoder.** These tasks are drawn from the Alexandria materials dataset [7].

A.12.3 Composition + NN

Three-layer feed-forward neural network: This model comprises three fully connected layers, each followed by a ReLU activation, optional dropout for regularization, and layer normalization to accelerate convergence and stabilize training. Compositional feature vectors (for example, elemental fractional embeddings) are input to this multilayer perceptron (MLP) to predict the target property.

A.12.4 Composition + XGB

A gradient-boosted decision-tree model trained on the same compositional features. XGBoost captures nonlinear interactions among elemental descriptors through an ensemble of shallow trees, providing an efficient and robust baseline for materials-property prediction [82].

A.13 LM Prompt Template

Listing 1 provides the complete prompt used to elicit precursor proposals for solid-state synthesis of $Na_3Bi(AsO_4)_2$. It directs the model to produce 20 candidate precursor sets, each covering all elements of the target, prioritizing stable and commonly used reagents, and ordering the sets by chemical plausibility.

```
You are a computational chemistry expert specializing in solid-state
    synthesis and retrosynthesis.
Your task is to identify potential precursor combinations for solid-
    state synthesizing the target material: 'Na3Bi(AsO4)2'.

**Requirements**:
1.  Generate 20 distinct combinations of precursor materials.
2.  Use standard chemical formulas ONLY (e.g., 'TiO2', 'Na2CO3').
3.  **Constraint Check:** Ensure each precursor combination contains
    ALL elements present in the target material 'Na3Bi(AsO4)2'. Assume
     Oxygen and other common laboratory elements (e.g., C for
    carbonate sources) are available.
4.  **Plausibility Filter:** Prefer chemically plausible routes using
    reasonably common and stable laboratory reagents. A plausible
    route is one that uses precursors commonly found in solid-state
    synthesis and avoids highly unstable or rare compounds.
5.  Order the 20 combinations from the MOST plausible/common synthesis
     routes to the LEAST plausible/common.
6.  **Common Precursor Types:** Consider oxides (e.g., TiO2, Fe2O3),
    carbonates (e.g., Na2CO3, CaCO3), nitrates (e.g., KNO3, Ca(NO3)2),
     hydroxides (e.g., Al(OH)3), and other standard laboratory
    reagents.
7.  **No Gases:** Do not include 'O2' in the precursor combinations.
7.  If the target material is not suitable for solid-state synthesis,
    respond with False as a boolean.

**Examples of Target -> Precursors:**
- Target: "\ce{MoF5}", Precursors: [False] #only synthesizable via gas
    -solid reaction, not suitable for conventional solid-state
    synthesis
- Target: 'Gd2TiO5', Precursors: ['TiO2', 'Gd2O3']
- Target: 'NdTl(MoO4)2', Precursors: ['MoO3', 'Tl2O3', 'Nd2O3']
- Target: 'Sr(GaO2)2', Precursors: ['SrCO3', 'Ga2O3']
- Target: 'La0.075Ta2O5.113', Precursors: ['La2O3', 'Ta2O5']
- Target: 'LaFeO3', Precursors: ['Fe2O3', 'LaCO3']
- Target: 'Sr1.9Ca1Tl0.9V0.1Cu2Bi0.1O7', Precursors: ['SrCO3', 'Tl2O3'
    , 'CuO', 'CaO', 'Bi2O3', 'V2O5']
- Target: 'La3RuO7', Precursors: ['RuO2', 'La2O3']
- Target: 'Zr0.8Ti1Sn0.2O4', Precursors: ['SnO2', 'TiO2', 'ZrO2']
- Target: 'La1Fe0.95W0.05O3', Precursors: ['WO3', 'Fe2O3', 'La2O3']
- Target: 'Bi3PO7', Precursors: ['PH9(NO2)2', 'Bi2O3']
- Target: 'CsTaWO6', Precursors: ['WO3', 'Cs2CO3', 'Ta2O5']
- Target: 'TaWO6', Precursors: ['WO3', 'Ta2O5']
- Target: 'Dy0.05Zn1Ga1.95O4', Precursors: ['Dy2O3', 'Ga2O3', 'ZnO']
```

```
- Target: 'Nd0.02Gd0.98V1O4', Precursors: ['Gd2O3', 'V2O5', 'Nd2O3']
- Target: 'Ba1Pr0.8In0.2O3', Precursors: ['BaCO3', 'Pr6O11', 'In2O3']
- Target: 'Ba4SrSmTi3V7O30', Precursors: ['SrCO3', 'Sm2O3', 'BaCO3', '
    TiO2', 'V2O5']
- Target: 'ZrSiO', Precursors: ['SiO2', 'ZrO2']
- Target: 'Ba0.6Sr0.4Nb0.1Co0.9O3', Precursors: ['SrCO3', 'Nb2O5', '
    BaCO3', 'Co2O3']
- Target: 'CsAlP2O7', Precursors: ['Cs2O', 'P2O5', 'Al2O3']
- Target: 'Ag4.64Pb2O5.87', Precursors: ['Ag2O', 'PbO']
- Target: 'Ca1Ti4Cu3.2O12', Precursors: ['CuO', 'TiO2', 'CaCO3']
- Target: 'Gd0.3Fe1Bi0.7O3', Precursors: ['Gd2O3', 'Bi2O3', 'Fe2O3']
- Target: 'YTiO', Precursors: ['TiO2', 'Y2O3']
- Target: 'Ba6Sn6Se13', Precursors: ['BaSe', 'Se', 'Sn']
- Target: 'V0.9Cu0.1Bi2O5.35', Precursors: ['CuO', 'Bi2O3', 'V2O5']
- Target: 'MgNb2(PbO3)3', Precursors: ['Nb2O5', 'MgCO3', 'PbO']
- Target: 'Cs0.75Rb0.25P1H2O4', Precursors: ['RbP(HO2)2', 'CsP(HO2)2']
- Target: 'TiCdO3', Precursors: ['TiO2', 'CdO']
- Target: 'Pu0.9Am0.1O2', Precursors: ['AmO2', 'PuO2']
- Target: 'La0.9Mn1Pb0.1O3', Precursors: ['PbO', 'MnO2', 'La2O3']
- Target: 'Li3.55Ca5.45Si3O12.45F1.55', Precursors: ['SiO2', 'CaCO3',
    'Li2CO3', 'LiF']
- Target: 'HoMnO3', Precursors: ['Mn2O3', 'Ho2O3']
- Target: 'SiPbC', Precursors: ['Pb', 'SiC']
- Target: 'CsAl(SiO3)2', Precursors: ['SiO2', 'Cs2CO3', 'Al2O3']
- Target: 'CdIn2O4', Precursors: ['CdO', 'In2O3']
- Target: 'CdWO4', Precursors: ['WO3', 'CdO']
- Target: 'Sr1.8Ca0.9Ti0.2Tl0.9Cu2Bi0.1O7', Precursors: ['SrCO3', '
    Tl2O3', 'Ti2O3', 'CuO', 'CaO', 'Bi2O3']
- Target: 'Ca3ZrSi2O9', Precursors: ['SiO2', 'CaCO3', 'ZrO2']
- Target: 'Ba3NbFe3(SiO7)2', Precursors: ['SiO2', 'Nb2O3', 'BaCO3', '
    Fe2O3']
- Target: 'Mn2Ni(CO2)6', Precursors: ['MnH6(CO)4', 'NiH6(CO)4']

Note: Ensure the quality and consistency of the example data to
    prevent parsing issues.

Generate the list of 20 precursor combinations for the target 'Na3Bi(
    AsO4)2'.
**Output Format:** Respond ONLY with a single Python-formatted list of
     lists. Each inner list should contain the precursor strings.
    Typically, 2-4 precursors per combination are expected.
Example Output Format: [['precursor1a', 'precursor1b', 'precursor1c'],
     ['precursor2a', 'precursor2b', 'precursor2c'], , ..., ['
    precursor20a', 'precursor20b', 'precursor20c']]

**Important:** Ensure all combinations are chemically valid and
    contain all elements needed to synthesize the target material. Do
    not include any explanations or text outside the Python list
    format.
```

Listing 1: Prompt for precursor generation

Conditions Prompt

```
\begin{lstlisting}
 You are a computational chemistry expert specializing in solid-state
     synthesis.
Assume there is only one sintering and one calcination step involved.
Your task is to predict the optimal synthesis conditions for the
    following chemical reaction:
4 BaCO3 + 1 Fe2O3 + 4 Nb2O5 + 0.333 Pr6O11 == 1 Ba4Pr2Fe2Nb8O30 + 4
    CO2 + 0.333 O2

**Required Conditions to Predict**:
1. Sintering Temperature (in \(^{\circ}\)C)
```

```
2. Sintering Time (in hours)
3. Calcination Temperature (in \(^{\circ}\)C)
4. Calcination Time (in hours)

**Guidelines for Prediction**:
- Base your predictions on established solid-state chemistry
    principles.
- Provide scientifically plausible values within typical laboratory
    ranges.
- Assume there is only one sintering and one calcination step involved
    .

**Examples of Synthesis Conditions:**
- Reaction: 2 BaCO3 + 0.667 Co3O4 + 6 Fe2O3 == 1 Ba2Co2Fe12O22 + 2 CO2
     + 0.333 O2
  Sintering Temperature (\(^{\circ}\)C): 1240.0
  Sintering Time (hours): 4.0
  Calcination Temperature (\(^{\circ}\)C): 1000.0
  Calcination Time (hours): 6.0

- Reaction: 1 GeO2 + 1 ZnO == 1 ZnGeO3
  Sintering Temperature (\(^{\circ}\)C): 1240.0
  Sintering Time (hours): 6.0
  Calcination Temperature (\(^{\circ}\)C): 1000.0
  Calcination Time (hours): 6.0

- Reaction: 0.5 In2O3 + 0.5 La2O3 == 1 LaInO3
  Sintering Temperature (\(^{\circ}\)C): 1150.0
  Sintering Time (hours): 24.0
  Calcination Temperature (\(^{\circ}\)C): 830.0
  Calcination Time (hours): 1.0

- Reaction: 0.05 Fe2O3 + 0.015 O2 + 0.98 SrCO3 + 0.92 TiO2 == 1 Sr0.98
    Ti0.92Fe0.1O3 + 0.98 CO2
  Sintering Temperature (\(^{\circ}\)C): 1450.0
  Sintering Time (hours): 24.0
  Calcination Temperature (\(^{\circ}\)C): 1100.0
  Calcination Time (hours): 12.0

- Reaction: 1.5 Li2CO3 + 0.5 Nb2O5 == 1 Li3NbO4 + 1.5 CO2
  Sintering Temperature (\(^{\circ}\)C): 970.0
  Sintering Time (hours): 2.0
  Calcination Temperature (\(^{\circ}\)C): 800.0
  Calcination Time (hours): 4.0

- Reaction: 0.05 Al2O3 + 0.9 CaCO3 + 0.05 Nd2O3 + 0.9 TiO2 == 1 Ca0.9
    Nd0.1Ti0.9Al0.1O3 + 0.9 CO2
  Sintering Temperature (\(^{\circ}\)C): 1300.0
  Sintering Time (hours): 12.0
  Calcination Temperature (\(^{\circ}\)C): 1200.0
  Calcination Time (hours): 10.0

- Reaction: 1 BaCO3 + 1 Ir + 1 O2 == 1 BaIrO3 + 1 CO2
  Sintering Temperature (\(^{\circ}\)C): 1000.0
  Sintering Time (hours): 72.0
  Calcination Temperature (\(^{\circ}\)C): 900.0
  Calcination Time (hours): 12.0

- Reaction: 2 BaCO3 + 3 CuO + 0.5 Y2O3 == 1 Y1Ba2Cu3O + 2 CO2 + 2.75
    O2
  Sintering Temperature (\(^{\circ}\)C): 950.0
  Sintering Time (hours): 8.0
  Calcination Temperature (\(^{\circ}\)C): 950.0
  Calcination Time (hours): 24.0
```

```
- Reaction: 1.5 Bi2O3 + 0.5 O2 + 0.5 Sb2O3 == 1 Bi3SbO7
  Sintering Temperature (\(^{\circ}\)C): 860.0
  Sintering Time (hours): 2.0
  Calcination Temperature (\(^{\circ}\)C): 700.0
  Calcination Time (hours): 4.0

- Reaction: 0.667 Co3O4 + 1 SnO2 == 1 Co2SnO4 + 0.333 O2
  Sintering Temperature (\(^{\circ}\)C): 1250.0
  Sintering Time (hours): 48.0
  Calcination Temperature (\(^{\circ}\)C): 1200.0
  Calcination Time (hours): 48.0

- Reaction: 0.5 O2 + 1 PbO + 2 TiO2 == 1 Pb(TiO3)2
  Sintering Temperature (\(^{\circ}\)C): 1100.0
  Sintering Time (hours): 2.0
  Calcination Temperature (\(^{\circ}\)C): 800.0
  Calcination Time (hours): 2.0

- Reaction: 0.8 CaCO3 + 0.2 SrCO3 + 1 TiO2 == 1 Ca0.8Sr0.2TiO3 + 1 CO2
  Sintering Temperature (\(^{\circ}\)C): 1460.0
  Sintering Time (hours): 4.0
  Calcination Temperature (\(^{\circ}\)C): 1100.0
  Calcination Time (hours): 4.0

- Reaction: 1 Li2CO3 + 1 TiO2 == 1 Li2TiO3 + 1 CO2
  Sintering Temperature (\(^{\circ}\)C): 1100.0
  Sintering Time (hours): 2.0
  Calcination Temperature (\(^{\circ}\)C): 600.0
  Calcination Time (hours): 2.0

- Reaction: 2 TiO2 + 1 Y2O3 == 1 Y2Ti2O7
  Sintering Temperature (\(^{\circ}\)C): 1600.0
  Sintering Time (hours): 6.0
  Calcination Temperature (\(^{\circ}\)C): 1000.0
  Calcination Time (hours): 6.0

- Reaction: 1 La2O3 + 2 ZrO2 == 1 La2Zr2O7
  Sintering Temperature (\(^{\circ}\)C): 1549.85
  Sintering Time (hours): 50.0
  Calcination Temperature (\(^{\circ}\)C): 1399.85
  Calcination Time (hours): 10.0

- Reaction: 0.5 Dy2O3 + 0.25 O2 + 1 RuO2 + 2 SrCO3 == 1 Sr2DyRuO6 + 2
    CO2
  Sintering Temperature (\(^{\circ}\)C): 1199.85
  Sintering Time (hours): 24.0
  Calcination Temperature (\(^{\circ}\)C): 979.85
  Calcination Time (hours): 12.0

- Reaction: 0.025 Gd2O3 + 0.013 O2 + 0.95 SnO2 == 1 Gd0.05Sn0.95O2
  Sintering Temperature (\(^{\circ}\)C): 1350.0
  Sintering Time (hours): 24.0
  Calcination Temperature (\(^{\circ}\)C): 1200.0
  Calcination Time (hours): 24.0

- Reaction: 1 MgO + 1 TiO2 == 1 MgTiO3
  Sintering Temperature (\(^{\circ}\)C): 1400.0
  Sintering Time (hours): 4.0
  Calcination Temperature (\(^{\circ}\)C): 1100.0
  Calcination Time (hours): 4.0

- Reaction: 1.7 CaCO3 + 2 CuO + 0.15 La2O3 + 4 TiO2 == 1 Ca1.7La0.3
    Cu2Ti4O12 + 1.7 CO2 + 0.075 O2
  Sintering Temperature (\(^{\circ}\)C): 1100.0
  Sintering Time (hours): 24.0
```

```
  Calcination Temperature (\(^{\circ}\)C): 900.0
  Calcination Time (hours): 5.0

- Reaction: 0.05 Fe2O3 + 0.035 O2 + 0.98 SrCO3 + 0.9 TiO2 == 1 Sr0.98
    Ti0.9Fe0.1O3 + 0.98 CO2
  Sintering Temperature (\(^{\circ}\)C): 1450.0
  Sintering Time (hours): 24.0
  Calcination Temperature (\(^{\circ}\)C): 1100.0
  Calcination Time (hours): 12.0

- Reaction: 1 MoO3 + 1 Nd2O3 == 1 Nd2MoO6
  Sintering Temperature (\(^{\circ}\)C): 1350.0
  Sintering Time (hours): 6.0
  Calcination Temperature (\(^{\circ}\)C): 1200.0
  Calcination Time (hours): 4.0

- Reaction: 1 Li2CO3 + 1 MnCO3 + 0.5 O2 == 1 Li2MnO3 + 2 CO2
  Sintering Temperature (\(^{\circ}\)C): 750.0
  Sintering Time (hours): 24.0
  Calcination Temperature (\(^{\circ}\)C): 700.0
  Calcination Time (hours): 20.0

- Reaction: 0.5 Fe2O3 + 0.5 Tb2O3 == 1 TbFeO3
  Sintering Temperature (\(^{\circ}\)C): 1300.0
  Sintering Time (hours): 12.0
  Calcination Temperature (\(^{\circ}\)C): 1200.0
  Calcination Time (hours): 12.0

- Reaction: 3 BaCO3 + 1.18 CaO + 0.91 Nb2O5 + 0.135 O2 == 1 Ba3Ca1.18
    Nb1.82O9 + 3 CO2
  Sintering Temperature (\(^{\circ}\)C): 1600.0
  Sintering Time (hours): 2.0
  Calcination Temperature (\(^{\circ}\)C): 1400.0
  Calcination Time (hours): 4.0

- Reaction: 1 CaCO3 + 3 CuO + 4 TiO2 == 1 CaCu3Ti4O12 + 1 CO2
  Sintering Temperature (\(^{\circ}\)C): 1080.0
  Sintering Time (hours): 10.0
  Calcination Temperature (\(^{\circ}\)C): 950.0
  Calcination Time (hours): 15.0

- Reaction: 1 CaCO3 + 1 ZrO2 == 1 CaZrO3 + 1 CO2
  Sintering Temperature (\(^{\circ}\)C): 1449.85
  Sintering Time (hours): 50.0
  Calcination Temperature (\(^{\circ}\)C): 1349.85
  Calcination Time (hours): 10.0

- Reaction: 0.9 CeO2 + 0.025 O2 + 0.05 Re2O3 == 1 Re0.1Ce0.9O2
  Sintering Temperature (\(^{\circ}\)C): 1650.0
```