

# PhysNLU: A Language Resource for Evaluating Natural Language Understanding and Explanation Coherence in Physics

Jordan Meadows<sup>1,2</sup>, Zili Zhou<sup>1</sup>, André Freitas<sup>1,2</sup>

<sup>1</sup>Department of Computer Science, University of Manchester

<sup>2</sup>Idiap Research Institute, Switzerland

jordan.meadows@postgrad.manchester.ac.uk

{zili.zhou, andre.freitas}@manchester.ac.uk

## Abstract

In order for language models to aid physics research, they must first encode representations of mathematical and natural language discourse which lead to coherent explanations, with correct ordering and relevance of statements. We present a collection of datasets developed to evaluate the performance of language models in this regard, which measure capabilities with respect to sentence ordering, position, section prediction, and discourse coherence. Analysis of the data reveals equations and sub-disciplines which are most common in physics discourse, as well as the sentence-level frequency of equations and expressions. We present baselines that demonstrate how contemporary language models are challenged by coherence related tasks in physics, even when trained on mathematical natural language objectives.

**Keywords:** mathematical text, physics, natural language understanding, discourse coherence

## 1. Introduction

Physics literature is a form of mathematical language which is unique beyond simply domain vocabulary. How physicists use mathematics to reason and explain, separates their field fundamentally from other disciplines, including mathematics. Many of its sub-disciplines are situated between pure mathematics and engineering, while others conjoin computer science and biology, with physical methods acting as a well-travelled bridge between the formal and natural sciences. It has not been proven, for example, that smooth solutions (Pizzocchero, 2021; Gala et al., 2021; Miller, 2021) always exist for the Navier-Stokes equations (a millenium problem) despite their widespread use in simulating and engineering fluid dynamics, while biophysics demonstrates that fundamental problems in ecology and evolution can be characterized by computational complexity classes (Ibsen-Jensen et al., 2015). Physics discourse serves as a universal mechanism for generating empirically falsifiable quantitative theory in the natural sciences and engineering (Smith and Fleck, 2017; Coffey and Kalmykov, 2012), separate to both pure mathematics and any downstream field.

Its core traits are reflected in unique literary devices and its mathematical explanations will differ to those in the formal sciences as a result. A concrete example is the physics derivation; a core explanatory or argumentative device less rigorous and more informal than mathematical proofs (Meadows and Freitas, 2021; Davis, 2019; Kaliszzyk et al., 2015), which generally results in predictive equations relating physical quantities, rather than generating a truth value for a given conjecture (*e.g.*, twin primes). Such equations are central components of physics descriptions, with natural language forming around them and their elements, and their relation to other equations through derivations.

Mathematics as a whole, particularly logic, is less concerned with this predictive modelling of real world systems, let alone when such systems are quantum or rela-

tivistic, or both. Suggesting that mathematicians work at a level of abstraction higher than that of physicists (*i.e.*, proof frameworks compared to specific derivations), Feynman famously states that “*Mathematicians are only dealing with the structure of reasoning...*”.

Within the unique sphere of physics literature, we introduce a suite of datasets which together gauge a model’s proficiency in recognising whether or not a physics-related explanation is coherent. In parallel with tasks inspired by DiscoEval (Chen et al., 2019), we aim to “evaluate the discourse-related knowledge captured by pretrained sentence representations” in the physics domain. From the proposed data we show that modern pretrained language models are challenged by these tasks even after fine-tuning, in particular demonstrating that a recent language model (Peng et al., 2021) with state-of-the-art performance in mathematical information retrieval, formula topic classification, and formula headline generation tasks, trained with equation-context pairs extracted from arXiv on math-aware pretraining objectives, is outperformed by even vanilla BERT-Base and all popular non-mathematical language models considered in this work.

We contribute the following:

1. We introduce PhysNLU; a collection of 4 core datasets related to sentence classification, ordering, and coherence of physics explanations based on related tasks (Chen et al., 2019). Each dataset comprises explanations extracted from Wikipedia including derivations and mathematical language. We additionally present 2 parent datasets extracted from 6.3k articles related to physics, in both raw Wikipedia data, and in a form that mimics WikiText-103 (Merity et al., 2016), which is a popular dataset used in related work (Iter et al., 2020). PhysNLU is available online<sup>1</sup>.

<sup>1</sup><https://github.com/jmeadows17/PhysNLU>

- We provide analysis of linguistic features of physics text, including insights such as sentence and example-level distribution of mathematical content across the datasets, the frequency in which explained concepts relate to physics sub-domains, and the most frequent equations in the discourse.
- We demonstrate how the state-of-the-art does not exhibit proficient inference capabilities with respect to tasks concerning order, coherence, relative position, and classification of sentence-level physics explanations, even when approaches have been designed explicitly for mathematical language, through baselines extracted from experiments involving a selection of pretrained language models.

## 2. Task Description

The tasks considered in this work probe model proficiency across 4 categories, originally designed for general language, but here employed specifically for physics discourse containing mathematics. Binary Sentence Ordering tests the ability of a model to recognise order at the shortest possible scale, between two sentences. Sentence Position tests this order and position recognition at a larger scale, closer to that of full paragraphs. Discourse Coherence tests whether a model can determine whether a sequence of statements in an explanation are continuous and relevant. Sentence Section Prediction tests how well a model can link individual sentences to a specific section of an explanation. Together, in our context, they evaluate the discourse-related knowledge captured by pretrained sentence representations, and physics explanation coherence with respect to order and sentence relevance. We now describe our method for data collection for each of the 6 datasets, including the 4 directly used in the forthcoming experiments for each task as described in Figure 2, while an overview of our contributions are displayed in Figure 1.

## 3. Dataset Collection

### 3.1. PhysNLU-WikiRaw

Starting from an English Wikipedia XML dump, we select articles with a mention of “physics” and contain at least one equation defined with a `<math>` tag. After cleaning articles to contain mostly mathematical natural language, and removing those which are predominantly tables, this results in a dataset containing 6.3k articles. We include article titles and corresponding raw unedited text, as well as wikipedia article categories.

### 3.2. PhysNLU-WikiText

This data mimics WikiText-103 (Merity et al., 2016) which is used for the approach introducing the CONPONO objective (Iter et al., 2020) during a preprocessing stage<sup>2</sup>. Among other similarities, we opt to nest section titles within equals signs (e.g. “ = Title = ”, “ = = Section = = ”)

<sup>2</sup>[https://github.com/google-research/language/blob/master/language/conpono/create\\_pretrain\\_data/wiki\\_preproc\\_pipeline.py](https://github.com/google-research/language/blob/master/language/conpono/create_pretrain_data/wiki_preproc_pipeline.py)

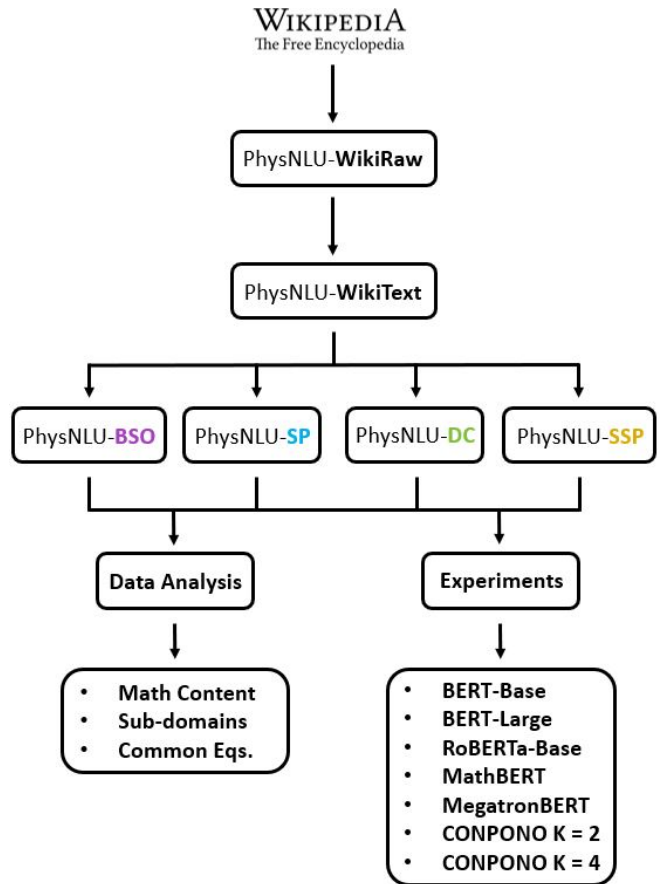


Figure 1: Workflow and overview.

and omit reference and “see more” sections. The major linguistic differences between WikiText-103 and PhysNLU-WikiText are the inclusion of mathematical content as well as structures which may contain mathematical expressions such as tables, which are infrequent. This core dataset is taken as the starting point from which to derive the other datasets. We then extract 516k unique sentences for use in the following datasets, where sentences are determined by splits on full stops which, similarly to commas, are separated by a space from words (e.g. “end of sentence .”). We correct for issues with names (e.g. J. J. Hopfield) and abbreviations, and some instances where full stops should be present but are omitted.

### 3.3. PhysNLU-BSO (Binary Sentence Ordering)

We take all pairs of consecutive sentences from sections, where selected pairs overlap. Each pair has a 50% chance that the pair order is swapped and we include a label to denote whether a swap (1) has occurred or not (0), suitable to be framed as binary classification. The BSO dataset contains 459k examples.

### 3.4. PhysNLU-SP (Sentence Position)

We take the first 5 sentences from each applicable section, select a sentence at random and move it to the first position (shifting the others down). The number of the swapped sentence is the label corresponding to each set of 5 sentences, suitable for multiclass classification. The SP dataset contains 40k examples.

### Sentence Position (SP)

This is the differential equations formulation of Gauss equation up to a trivial rearrangement. 4

According to the (purely mathematical) Gauss divergence theorem, the electric flux through the boundary surface  $\partial\Omega$  can be rewritten as

$$\oiint_{\partial\Omega} \mathbf{E} \cdot d\mathbf{S} = \iiint_{\Omega} \nabla \cdot \mathbf{E} dV \quad 1$$

The integral version of Gauss's equation can thus be rewritten as

$$\iiint_{\Omega} \left( \nabla \cdot \mathbf{E} - \frac{\rho}{\epsilon_0} \right) dV = 0 \quad 2$$

Since  $\Omega$  is arbitrary (e.g. an arbitrary small ball with arbitrary center), this is satisfied if and only if the integrand is zero everywhere. 3

### Discourse Coherence (DC)

According to the (purely mathematical) Gauss divergence theorem, the electric flux through the boundary surface  $\partial\Omega$  can be rewritten as

$$\oiint_{\partial\Omega} \mathbf{E} \cdot d\mathbf{S} = \iiint_{\Omega} \nabla \cdot \mathbf{E} dV \quad 1$$

The integral version of Gauss's equation can thus be rewritten as

$$\iiint_{\Omega} \left( \nabla \cdot \mathbf{E} - \frac{\rho}{\epsilon_0} \right) dV = 0 \quad 2$$

For that reason, it is called the heat equation in mathematics, even though it applies to many other physical quantities besides temperature. 3

This is the differential equations formulation of Gauss equation up to a trivial rearrangement. 4

### Binary Sentence Ordering (BSO)

The integral version of Gauss's equation can thus be rewritten as

$$\iiint_{\Omega} \left( \nabla \cdot \mathbf{E} - \frac{\rho}{\epsilon_0} \right) dV = 0 \quad 2$$

According to the (purely mathematical) Gauss divergence theorem, the electric flux through the boundary surface  $\partial\Omega$  can be rewritten as

$$\oiint_{\partial\Omega} \mathbf{E} \cdot d\mathbf{S} = \iiint_{\Omega} \nabla \cdot \mathbf{E} dV \quad 1$$

### Sentence Section Prediction (SSP)

According to the (purely mathematical) Gauss divergence theorem, the electric flux through the boundary surface  $\partial\Omega$  can be rewritten as

$$\oiint_{\partial\Omega} \mathbf{E} \cdot d\mathbf{S} = \iiint_{\Omega} \nabla \cdot \mathbf{E} dV$$

Figure 2: Each numbered box is a physics statement. Explanation of how to obtain the differential form of Gauss' law via the divergence theorem, used to demonstrate how 4 evaluation tasks handle physics explanations (extracted from Wikipedia, including errors). **(SP)** Sentence position (top left) takes a random sentence from the description and moves it to position 1, then the model predicts the true position of the new first sentence, which in this case is 4. **(DC)** Discourse coherence (top right) randomly replaces a sentence with another from a different physics explanation, where the model predicts if the explanation is coherent. **(BSO)**: Binary sentence ordering (bottom left) swaps two consecutive sentences, and the model predicts if the second entails the first. **(SSP)**: Sentence section prediction takes a random sentence, and the model predicts if it belongs to an introduction or otherwise.

### 3.5. PhysNLU-DC (Discourse Coherence)

The first 6 sentences from each applicable section are selected, then between positions 2 and 5 inclusive a sentence is swapped with another article at random, with 50% swap occurrence. Whether a swap has occurred or not is included as a label for each example for binary classification, and the DC dataset contains 35k examples.

### 3.6. PhysNLU-SSP (Sentence Section Prediction)

All sentences from the introduction sections of each article are selected and an equal number of sentences are extracted from elsewhere at random from the corpus. Introduction sentences are associated with a label (1) while non-introductory sentences are associated with a separate label (0) for binary classification. The SSP dataset contains 90k examples.

## 4. Dataset Statistics

We now analyse our data with a focus on equations and mathematical natural language. Table 1 shows an overview of notable features, such as the proportion of examples in each dataset which contains mathematical expressions, or specifically equations.

Figure 3 describes the proportion of sentences which contain at least  $n$  mathematical elements for  $n \in [1, 6]$ , where an element is identified via 3 separate tags:  $\langle \text{math} \rangle$ ,  $\{ \text{math} \}$ , and  $\{ \text{mvar} \}$ . The lighter bars correspond to all sentences present in the evaluation data, while the darker bars correspond to sentences in the SSP dataset which contain proportionally less math. Introductory sentences make up half of the data for SSP and usually they do not contain mathematical language or equations, which accounts for this gap. The proportion of math in sentences from the

BSO, SP, and DC datasets are practically equivalent to the overall proportion. Figure 4 shows the relative frequency and proportion of the top 8 Wikipedia categories associated with each article. A single article can correspond to a large number of categories, out of 12.5k categories in our case. Notably, fields related to quantum mechanics are by far the most frequent, where 10% of the data corresponds to either “Quantum mechanics”, “Quantum field theory”, or “Condensed matter physics”.

Figure 5 displays how often specific equations are present in the corpora. One might be tempted to claim that this demonstrates how physicists tend to argue and explain using initial conditions with respect to time ( $t = 0$ ), displacement ( $x = 0, r = 0, z = 0$ ), and angle ( $\theta = 0$ ), however this exact string matching is biased towards simple equations. As the complexity of equations increases to include multiple terms, and many terms are equivalent in meaning but different in notation, there will be multiple equations in the data which correspond to the same physics. A more accurate way to assess this would involve classifying groups of equations with a good math retrieval model (Peng et al., 2021) and counting group frequency. This analysis does offer insight for simple equations however. For example, it reflects the convention that people prefer to start counting from  $n = 1$  in physics, which occurs more frequently than  $n = 0$ , that the famous  $E = mc^2$  is more prolific than the similarly famous  $F = ma$ , and that the most frequently discussed Maxwell equation is  $\nabla \cdot \mathbf{B} = 0$ .

Figure 6 shows the proportion of examples from each evaluation dataset which contain at least  $n$  counts of either a  $\langle \text{math} \rangle$  equation or non-equational math, for  $n \in [1, 6]$ . The darker bars represent solely equations. DC, SP, BSO and SSP examples comprise 6, 5, 2, and 1 sentences respectively, so we expect that the proportion of math included decreases in that order.

Dataset	Size	% with math	% with equations
DC	35 k	45	35
SP	40 k	36	29
BSO	459 k	24	17
SSP	90 k	12	7

Table 1: Metadata for each evaluation dataset. The percentage of examples with math refers to the percentage of examples with LaTeX text between the XML tags  $\langle \text{math} \rangle$  and  $\langle / \text{math} \rangle$ . The percentage of equations refers to only math which contains at least one equality.

## 5. Results

We evaluate models on 4 tasks from the DiscoEval suite (Chen et al., 2019). We remove the PDTB and RST related tasks due to the lack of a linguistic framework for describing discourse relations in the physics context. Table 1 gives additional information regarding the data used for each task.

### 5.1. Evaluation Tasks

DiscoEval is “designed to evaluate discourse-related knowledge in pretrained sentence representations”. We

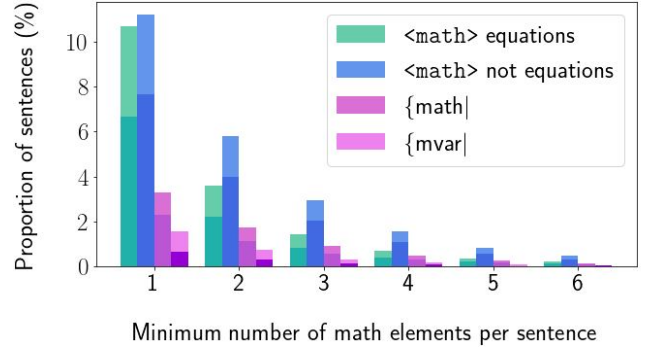


Figure 3: For 516k unique sentences used in the evaluation tasks, the proportion of sentences which contain at least  $n$  math elements is shown for  $n \in [1, 6]$ . The two  $\langle \text{math} \rangle$  variants correspond to LaTeX text identified by the XML  $\langle \text{math} \rangle$  tag which respectively do and do not contain an equality.  $\{ \text{math} \}$  and  $\{ \text{mvar} \}$  correspond to any mathematical text identified by each marker. The darker bars correspond to only sentences extracted from the SSP dataset.

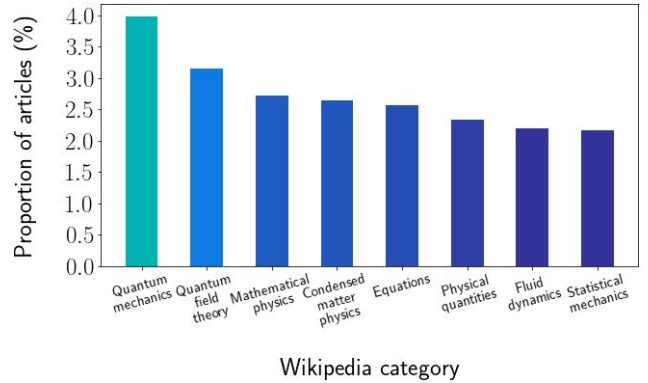


Figure 4: Top 8 most frequent article categories out of 12.5k, excluding categories containing the phrase “Articles containing ...”, where each article may correspond to multiple categories.

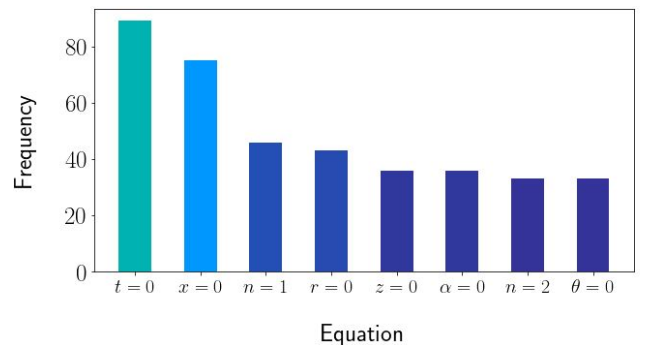


Figure 5: Top 8 most frequent  $\langle \text{math} \rangle$  tagged equations by exact string matching after accounting for spaces, commas, and full stops.

briefly describe the 4 evaluation tasks from DiscoEval considered in our work, with examples shown in Figure 2. We use the same conventions (Chen et al., 2019) for representing concatenation of vectors  $[., ., \dots]$ .

Sentence position (SP) involves considering 5 consecutive

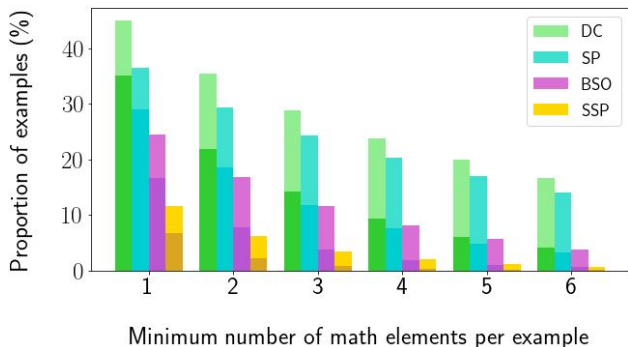


Figure 6: The proportion of examples which contain at least  $n$  math elements is shown for  $n \in [1, 6]$ , where examples are sourced from the DC, SP, BSO, and SSP datasets. The math is identified via the `<math>` tag, where lighter bars correspond to at least  $n$  math of any kind, while the darker bars correspond to the inclusion of only equations.

sentences at a time, moving a random sentence to the first position, then predicting the correct position of the first sentence. We take the first 5 sentences from every paragraph in our data. Classifiers are trained by encoding the 5 sentences to vector representations  $x_i$ , then vectors  $x_1 - x_i$  are concatenated to  $x_1$  for  $i \in [2, 5]$  as input to the classifier as:  $[x_1, x_1 - x_2, x_1 - x_3, x_1 - x_4, x_1 - x_5]$ .

Binary sentence ordering (BSO) involves taking pairs of contiguous sentences from a paragraph, swapping the order 50% of the time and predicting if a swap has occurred. A classifier is trained by concatenating  $x_1$  and  $x_2$  with their element-wise difference as:  $[x_1, x_2, x_1 - x_2]$ .

Discourse coherence (DC) involves taking 6 consecutive sentences, replacing a sentence from positions 2-5 inclusive with a sentence from a random article with 50% frequency and predicting if a swap has occurred. We take the first 6 sentences from each paragraph for this task. Each vector  $x_i$  is concatenated for input to the classifier as:  $[x_1, x_2, x_3, x_4, x_5, x_6]$ .

Sentence section prediction (SSP) involves sampling a sentence from either the abstract of a scientific article or elsewhere with equal probability, and predicting if the sentence belongs to the abstract. In our case, we sample from article introductions as we do not have abstracts. The original task (Chen et al., 2019) involves the omission of equations which increases the difficulty of the task, but due to the nature of our problem space we leave them in. The classifier input is just the vector representation  $x_1$ .

## 5.2. Baselines

We include 7 baseline transformer-architecture models (Vaswani et al., 2017) in this study, BERT-base-uncased, BERT-large-uncased (Devlin et al., 2019), RoBERTa-base (Liu et al., 2019), MathBERT (Peng et al., 2021), MegatronBERT (Shoeybi et al., 2019), CONPONO (K=2), CONPONO (K=4) (Iter et al., 2020).

BERT-base-uncased, BERT-large-uncased are two BERT models pretrained on large-scale common domain text corpora. Based on the 12 or 24 encoder layers, the models achieved state-of-the-art performances on several NLP tasks such as sentence classification, next sentence pre-

dition, token classification, etc. BERT-large-uncased has a larger parameter size than BERT-base-uncased. The smaller model outperforms the larger in our case.

RoBERTa-base has the same architecture as BERT but is pretrained with more data, different hyperparameters, and only full-length sentences, and shows the BERT models were undertrained.

The MathBERT model uses the BERT architecture, however equation-context pairs are jointly considered through multiple pretraining objectives, which also relies on extracted tree tuples (Davila et al., 2016; Davila and Zanibbi, 2017).

MegatronBERT model improves the architecture of the original BERT models to enable model deployment across distributed GPU environments, while simultaneously improving the accuracy of the model.

The CONPONO models use the encoder architecture (and same data) from BERT to encode text segments, but are pretrained instead with the CONPONO objective together with Masked Language Modelling (MLM). The purpose is to let the models learn discourse relationships between sentences with respect to order and distance, while using negatives to increase sentence representation quality. We use 2 versions of CONPONO as baselines in this paper, K=2 means considering a maximum of 2 sentences before or after the anchor segment during pretraining, K=4 means considering a maximum of 4 sentences before or after.

70-30 split testing and 5-fold split testing are conducted for each baseline. We list the results in Tables 2-5 and observe the following:

- MathBERT is outperformed by all other models in each task, including base and large vanilla BERT. Given the SOTA performance of the model on 3 tasks related to mathematical text (Peng et al., 2021), and MathBERT being the only model to train on equation-context pairs, this is a surprising result.
- BERT-base-uncased model outperforms BERT-large-uncased in each task.
- Comparing CONPONO K=2 and CONPONO K=4, the performances are similar in each task, with K=2 being marginally better.
- MegatronBERT outperforms all models in all tasks except for discourse coherence (DC), but its parameter size is larger than most of other baselines including BERT-based-uncased and CONPONO models.
- CONPONO K=2 model outperforms BERT-base-uncased, RoBERTa-base, and MathBERT on SP and BSO tasks.
- All baselines perform poorly on the DC task.
- All baselines perform relatively well on the SSP task, but the accuracy performances are close, with no massive margin even between all-around worst performer MathBERT and generally good MegatronBERT.

	70-30 Split				K-fold			
	Acc	F1	AP	ROC	Acc	F1	AP	ROC
BERT-base-uncased	0.332	0.320	0.346	0.676	0.340±0.010	0.314±0.007	0.349±0.004	0.678±0.004
BERT-large-uncased	0.306	0.254	0.318	0.651	0.306±0.006	0.269±0.017	0.317±0.004	0.651±0.005
RoBERTa-base	0.281	0.227	0.283	0.611	0.287±0.008	0.217±0.011	0.291±0.003	0.615±0.003
MathBERT	0.258	0.248	0.256	0.571	0.259±0.003	0.245±0.005	0.257±0.003	0.574±0.004
MegatronBERT	0.647	0.646	0.729	0.893	0.647±0.003	0.646±0.003	0.729±0.002	0.894±0.001
CONPONO K=2	0.536	0.536	0.577	0.820	0.532±0.005	0.528±0.005	0.574±0.006	0.816±0.003
CONPONO K=4	0.527	0.523	0.571	0.816	0.526±0.008	0.522±0.009	0.572±0.009	0.815±0.004

Table 2: Sentence Position

	70-30 Split				K-fold			
	Acc	F1	AP	ROC	Acc	F1	AP	ROC
BERT-base-uncased	0.535	0.530	0.543	0.550	0.517±0.010	0.444±0.067	0.539±0.001	0.546±0.001
BERT-large-uncased	0.521	0.477	0.529	0.532	0.517±0.010	0.467±0.069	0.528±0.001	0.532±0.001
RoBERTa-base	0.531	0.529	0.537	0.546	0.531±0.002	0.525±0.007	0.538±0.001	0.546±0.001
MathBERT	0.511	0.443	0.519	0.518	0.513±0.001	0.492±0.023	0.520±0.002	0.520±0.002
MegatronBERT	0.762	0.762	0.855	0.853	0.762±0.001	0.762±0.001	0.855±0.001	0.853±0.001
CONPONO K=2	0.671	0.671	0.741	0.742	0.670±0.002	0.670±0.003	0.741±0.002	0.741±0.002
CONPONO K=4	0.665	0.664	0.734	0.735	0.666±0.001	0.665±0.001	0.734±0.001	0.735±0.001

Table 3: Binary Sentence Ordering

	70-30 Split				K-fold			
	Acc	F1	AP	ROC	Acc	F1	AP	ROC
BERT-base-uncased	0.529	0.527	0.534	0.541	0.522±0.011	0.491±0.044	0.535±0.004	0.539±0.004
BERT-large-uncased	0.505	0.402	0.519	0.522	0.514±0.006	0.435±0.056	0.520±0.007	0.524±0.008
RoBERTa-base	0.517	0.498	0.529	0.539	0.515±0.012	0.447±0.044	0.533±0.004	0.541±0.005
MathBERT	0.507	0.501	0.506	0.506	0.506±0.004	0.494±0.015	0.509±0.004	0.512±0.005
MegatronBERT	0.529	0.515	0.539	0.546	0.524±0.003	0.514±0.012	0.534±0.002	0.541±0.003
CONPONO K=2	0.534	0.532	0.538	0.545	0.533±0.003	0.532±0.004	0.536±0.003	0.544±0.004
CONPONO K=4	0.531	0.531	0.539	0.545	0.529±0.004	0.524±0.005	0.539±0.005	0.545±0.005

Table 4: Discourse Coherence

	70-30 Split				K-fold			
	Acc	F1	AP	ROC	Acc	F1	AP	ROC
BERT-base-uncased	0.695	0.694	0.756	0.763	0.693±0.003	0.693±0.003	0.754±0.002	0.762±0.002
BERT-large-uncased	0.670	0.670	0.723	0.733	0.666±0.003	0.665±0.003	0.718±0.002	0.728±0.002
RoBERTa-base	0.649	0.648	0.706	0.713	0.651±0.002	0.650±0.002	0.709±0.002	0.715±0.002
MathBERT	0.618	0.618	0.679	0.677	0.619±0.006	0.619±0.006	0.677±0.003	0.676±0.004
MegatronBERT	0.705	0.705	0.778	0.782	0.704±0.003	0.703±0.003	0.776±0.001	0.779±0.002
CONPONO K=2	0.687	0.686	0.761	0.756	0.685±0.004	0.684±0.004	0.746±0.002	0.754±0.003
CONPONO K=4	0.684	0.683	0.758	0.755	0.683±0.003	0.683±0.003	0.743±0.002	0.752±0.002

Table 5: Sentence Section Prediction

## 6. Related Work

DiscoEval (Chen et al., 2019) is a suite of evaluation tasks with the purpose of determining whether sentence representations include information about the role of a sentence in its discourse context. They build sentence encoders capable of modelling discourse information via training objectives that make use of natural annotations from Wikipedia, such as nesting level, section and article titles, among others.

Other core work (Iter et al., 2020) involves pretraining on both MLM and a contrastive inter-sentence objective (CONPONO), where they achieve state-of-the-art benchmarks for five of seven tasks in the DiscoEval suite, outperforming BERT-Large despite equalling the size of BERT-Base and training on the same amount of data. BERT-Base pretrained additionally on BSO in place of CONPONO, and BERT-large, claim the remaining two benchmarks.

Using full encoder-decoder transformer (Vaswani et al., 2017) architecture and an additional masked attention map which incorporates relationships between nodes in operator trees (OPTs) of equations (Davila et al., 2016; Davila and Zanibbi, 2017), MathBERT (Peng et al., 2021) approach pretrains with three objectives on arXiv data each extracting a specific latent aspect of information. MLM learns text representations, context correspondence prediction learns the latent relationship between formula and context, and masked substructure prediction learns semantic-level structure of formulas by predicting parent and child nodes in OPTs. MathBERT is fine-tuned and evaluated on mathematical information retrieval, formula topic classification, and formula headline generation, and outperforms previous approaches in each task. They train on 8.7 million equation-context pairs extracted from arXiv.

A data extraction pipeline (Ferreira and Freitas, 2020) collects 20k entries related to mathematical proofs from the ProofWiki website<sup>3</sup>, such as definitions, lemmas, corollaries, and theorems. They evaluate BERT and SciBERT by fine-tuning on a pairwise relevance classification task with their NL-PS dataset, where they classify if one mathematical text is related to another. As we highlight in the Introduction, physics and mathematics literature differ in their overarching considerations, and more specifically, the unstructured informal physics Wikipedia explanations that we present in our data naturally differ from the structured proofs present in NL-PS.

Their work builds on previous efforts applying NLP to general mathematics. One early approach (Zinn, 2003) proposes proof representation structures via discourse representation theory, including a prototype for generating formal proofs from informal mathematical discourse. Another approach (Cramer et al., 2009) focuses on development of a controlled natural language for mathematical texts which is compatible with existing proof verification software. Since these early developments, natural language-based problem solving and theorem proving have progressed significantly, but are still below human-level performance. For example, efforts towards building datasets for evaluating math word problem solvers (Huang et al., 2016) concludes the task as a significant challenge, with more recent large-scale dataset

construction and evaluation work (Amini et al., 2019; Miao et al., 2021) confirming that model performance is still well below the gold-standard. This difficulty extends to approaches involving pre-university math problems and geometric quantities (Matsuzaki et al., 2017; Lu et al., 2021). For automated theorem proving and mathematical reasoning, various datasets and accompanying approaches have been proposed (Kaliszyk et al., 2017; Bansal et al., 2019) including more recent work with equational logic (Piepenbrock et al., 2021) and language models (Rabe et al., 2020; Han et al., 2021).

A dataset construction approach and accompanying heuristic search for automating small physics derivations has been developed (Meadows and Freitas, 2021) which allows published results in modern physics to be converted into a form interpretable by a computer algebra system (Meurer et al., 2017), which then accommodates limited informal mathematical exploration. Detailed physics derivation data is scarce, and others have tackled such issues via synthetic data (Aygün et al., 2020), though not yet in physics. Reinforcement learning has been employed (Luo and Liu, 2018) to solve differential equations in nuclear physics with a template mapping method, and proof discovery and verification has been explored in relativity (Govindarajalulu et al., 2015). Theorem proving and derivation automation in physics remains elusive, with detailed discussions available in the literature (Kaliszyk et al., 2015; Davis, 2019).

With language models demonstrating logical capabilities with respect to type inference, missing assumption suggestion, and completing equalities (Rabe et al., 2020), as well as state-of-the-art performance in math retrieval and tasks related to equation-context correspondence (Peng et al., 2021), we believe our present work will contribute towards physics natural language / equational reasoners capable of generating coherent mathematical explanations and derivations.

## 7. Conclusion

Within the domain of physics, we present 2 parent datasets for general use and 4 specific datasets corresponding to discourse evaluation tasks (Chen et al., 2019) collectively referred as PhysNLU. The presented data frequently features equations, formulae, and mathematical language. Our analysis reveals that concepts related to quantum mechanics are most commonly discussed as determined by Wikipedia article category, that equations related to initial and boundary conditions are the most frequently considered when considering near-exact string matching, and we report the proportion of sentences and dataset examples which contain equations and mathematical terms identified by annotation frameworks native to Wikipedia. Finally, we present baseline results for popular non-mathematical language models and demonstrate that, despite expensive pretraining efforts and specialised training objectives for learning various aspects of mathematical text, such efforts do not improve the performance of language models in tasks related to sentence ordering, position, and recognising whether physics explanations are coherent. Future work will involve developing objectives which aid performance in this regard.

<sup>3</sup><https://proofwiki.org/wiki/MainPage>

## Acknowledgements

This work was partially funded by the SNSF project NeuMath (200021\_204617).

## 8. Bibliography

- Amini, A., Gabriel, S., Lin, P., Koncel-Kedziorski, R., Choi, Y., and Hajishirzi, H. (2019). Mathqa: Towards interpretable math word problem solving with operation-based formalisms. *arXiv preprint arXiv:1905.13319*.
- Aygün, E., Ahmed, Z., Anand, A., Firoiu, V., Glorot, X., Orseau, L., Precup, D., and Mourad, S. (2020). Learning to prove from synthetic theorems. *arXiv preprint arXiv:2006.11259*.
- Bansal, K., Loos, S., Rabe, M., Szegedy, C., and Wilcox, S. (2019). Holist: An environment for machine learning of higher order logic theorem proving. In *International Conference on Machine Learning*, pages 454–463. PMLR.
- Chen, M., Chu, Z., and Gimpel, K. (2019). Evaluation benchmarks and learning criteria for discourse-aware sentence representations. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 649–662, Hong Kong, China, November. Association for Computational Linguistics.
- Coffey, W. and Kalmykov, Y. P. (2012). *The Langevin equation: with applications to stochastic problems in physics, chemistry and electrical engineering*, volume 27. World Scientific.
- Cramer, M., Fisseni, B., Koepke, P., Kühlwein, D., Schröder, B., and Veldman, J. (2009). The naproche project controlled natural language proof checking of mathematical texts. In *International Workshop on Controlled Natural Language*, pages 170–186. Springer.
- Davila, K. and Zanibbi, R. (2017). Layout and semantics: Combining representations for mathematical formula search. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1165–1168.
- Davila, K., Zanibbi, R., Kane, A., and Tompa, F. W. (2016). Tangent-3 at the ntcir-12 mathir task. In *NTCIR*.
- Davis, E. (2019). Proof verification technology and elementary physics. In *Algorithms and Complexity in Mathematics, Epistemology, and Science*, pages 81–132. Springer.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019). Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*.
- Ferreira, D. and Freitas, A. (2020). Natural language premise selection: Finding supporting statements for mathematical text.
- Gala, S., Galakhov, E., Ragusa, M. A., and Salieva, O. (2021). Beale–kato–majda regularity criterion of smooth solutions for the hall–mhd equations with zero viscosity. *Bulletin of the Brazilian Mathematical Society, New Series*, pages 1–13.
- Govindarajalulu, N. S., Bringsjord, S., and Taylor, J. (2015). Proof verification and proof discovery for relativity. *Synthese*, 192(7):2077–2094.
- Han, J. M., Rute, J., Wu, Y., Ayers, E. W., and Polu, S. (2021). Proof artifact co-training for theorem proving with language models. *arXiv preprint arXiv:2102.06203*.
- Huang, D., Shi, S., Lin, C.-Y., Yin, J., and Ma, W.-Y. (2016). How well do computers solve math word problems? large-scale dataset construction and evaluation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 887–896.
- Ibsen-Jensen, R., Chatterjee, K., and Nowak, M. A. (2015). Computational complexity of ecological and evolutionary spatial dynamics. *Proceedings of the National Academy of Sciences*, 112(51):15636–15641.
- Iter, D., Guu, K., Lansing, L., and Jurafsky, D. (2020). Pre-training with contrastive sentence objectives improves discourse performance of language models.
- Kaliszyk, C., Urban, J., Siddique, U., Khan-Afshar, S., Dunchev, C., and Tahar, S. (2015). Formalizing physics: automation, presentation and foundation issues. In *International Conference on Intelligent Computer Mathematics*, pages 288–295. Springer.
- Kaliszyk, C., Chollet, F., and Szegedy, C. (2017). Holstep: A machine learning dataset for higher-order logic theorem proving. *arXiv preprint arXiv:1703.00426*.
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., and Stoyanov, V. (2019). Roberta: A robustly optimized bert pretraining approach. *ArXiv*, abs/1907.11692.
- Lu, P., Gong, R., Jiang, S., Qiu, L., Huang, S., Liang, X., and Zhu, S.-C. (2021). Theorem-aware geometry problem solving with symbolic reasoning and theorem prediction.
- Luo, M. and Liu, L. (2018). Automatic derivation of formulas using reinforcement learning. *arXiv preprint arXiv:1808.04946*.
- Matsuzaki, T., Ito, T., Iwane, H., Anai, H., and Arai, N. H. (2017). Semantic parsing of pre-university math problems. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2131–2141.
- Meadows, J. and Freitas, A. (2021). Similarity-based equational inference in physics. *Physical Review Research*, 3(4), Oct.
- Merity, S., Xiong, C., Bradbury, J., and Socher, R. (2016). Pointer sentinel mixture models. *arXiv preprint arXiv:1609.07843*.
- Meurer, A., Smith, C. P., Paprocki, M., Čertík, O., Kirpichev, S. B., Rocklin, M., Kumar, A., Ivanov, S., Moore, J. K., Singh, S., et al. (2017). Sympy: symbolic computing in python. *PeerJ Computer Science*, 3:e103.
- Miao, S.-Y., Liang, C.-C., and Su, K.-Y. (2021). A diverse corpus for evaluating and developing english math word problem solvers. *arXiv preprint arXiv:2106.15772*.
- Miller, E. (2021). A survey of geometric constraints on the blowup of solutions of the navier–stokes equation. *Journal of Elliptic and Parabolic Equations*, pages 1–11.



- Peng, S., Yuan, K., Gao, L., and Tang, Z. (2021). Mathbert: A pre-trained model for mathematical formula understanding.
- Piepenbrock, J., Heskes, T., Janota, M., and Urban, J. (2021). Learning equational theorem proving. *arXiv preprint arXiv:2102.05547*.
- Pizzocchero, L. (2021). On the global stability of smooth solutions of the navier–stokes equations. *Applied Mathematics Letters*, 115:106970.
- Rabe, M. N., Lee, D., Bansal, K., and Szegedy, C. (2020). Mathematical reasoning via self-supervised skip-tree training. *arXiv preprint arXiv:2006.04757*.
- Shoeybi, M., Patwary, M. A., Puri, R., LeGresley, P., Casper, J., and Catanzaro, B. (2019). Megatron-lm: Training multi-billion parameter language models using model parallelism. *ArXiv*, abs/1909.08053.
- Smith, R. W. and Fleck, C. (2017). Derivation and use of mathematical models in systems biology. In *Pollen Tip Growth*, pages 339–367. Springer.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. (2017). Attention is all you need.
- Zinn, C. (2003). A computational framework for understanding mathematical discourse. *Logic Journal of IGPL*, 11(4):457–484.