

MuLMS: A Multi-Layer Annotated Text Corpus for Information Extraction in the Materials Science Domain

Timo Pierre Schrader^{1,3} Matteo Finco² Stefan Grünewald^{1,4}
Felix Hildebrand² Annemarie Friedrich³

¹Bosch Center for Artificial Intelligence, Renningen, Germany

²Robert Bosch GmbH, Stuttgart, Germany

³University of Augsburg, Germany ⁴University of Stuttgart, Germany

timo.schrader|matteo.finco|stefan.gruenewald@de.bosch.com

annemarie.friedrich@informatik.uni-augsburg.de

Abstract

Keeping track of all relevant recent publications and experimental results for a research area is a challenging task. Prior work has demonstrated the efficacy of information extraction models in various scientific areas. Recently, several datasets have been released for the yet understudied materials science domain. However, these datasets focus on sub-problems such as parsing synthesis procedures or on sub-domains, e.g., solid oxide fuel cells.

In this resource paper, we present MuLMS, a new dataset of 50 open-access articles, spanning seven sub-domains of materials science. The corpus has been annotated by domain experts with several layers ranging from named entities over relations to frame structures. We present competitive neural models for all tasks and demonstrate that multi-task training with existing related resources leads to benefits.

1 Introduction

Designing meaningful experiments in empirical sciences requires maintaining a detailed overview of the huge amounts of literature published every year. Applying natural language processing (NLP) in this context has risen to be an active research area (Chandrasekaran et al., 2020; Beltagy et al., 2021; Cohan et al., 2022). Besides the biomedical field, which has been studied extensively in the past decades (e.g., Collier et al., 2004; Cohen et al., 2012; Demner-Fushman et al., 2022), the less-studied materials science domain has recently received more attention (Mysore et al., 2019; Friedrich et al., 2020; O’Gorman et al., 2021).

Materials science research aims to design and discover new materials. Part of the papers is hence often dedicated to the *synthesis procedures*, the “recipe” for creating a material. Their extraction from papers has been covered by Mysore et al. (2019) and O’Gorman et al. (2021). Much materials science research develops materials in the context of creating a particular device, e.g., batteries or

photovoltaic panels. The device is tested in various conditions and the literature needs to be analysed for identifying promising set-ups. Friedrich et al. (2020) address this for solid oxide fuel cells.

In this paper, we introduce **MuLMS** (the Multi-Layer Materials Science corpus), a new dataset of scientific publications annotated by domain experts with named entity (NE) mentions, relations, and frame structures corresponding to a broad notion of measurements (see Figure 1). In contrast to prior works, we include papers from a variety of materials science sub-domains. To the best of our knowledge, the existing datasets only annotate particular paragraphs or subsets of the sentences with NE mentions. Our dataset is the first to *exhaustively* annotate a large-scale collection of materials science articles with NEs and facilitates novel semantic search applications, e.g., answering search queries such as “find a passage within a paper reporting a measurement using material X, condition Y, and obtaining a value of at least Z.”

The design of MuLMS’ annotation scheme results from a collaboration of NLP and materials science experts. Our inter-annotator agreement study shows good agreement for most categories and decisions. We propose several machine learning tasks on the annotated data and present strong neural baselines for all tasks, which signals a high level of consistency across the annotations in our dataset. We cast detecting sentences describing measurements as a sentence classification task and provide a robust tagger for recognizing NEs. We propose to treat relation and semantic role extraction on MuLMS in a single step using a dependency parser that predicts relations between the NEs in a sentence. According to our multi-task experiments with related datasets, training jointly with MuLMS is beneficial for performance on those datasets.

Our contributions are as follows. (1) We publicly release a dataset of 50 open-access scientific publications *exhaustively* annotated by a domain

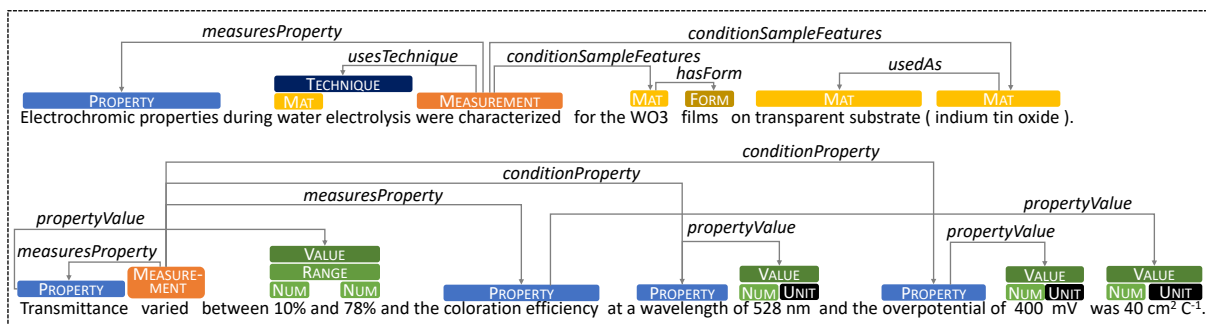


Figure 1: **Multi-Layer Materials Science Corpus**: named entity, relation and semantic role annotations.

expert with NE mentions, relations, and measurement frames.¹ (2) We define a set of NLP tasks on MuLMS and provide strong transformer-based baselines. Our code will be published. (3) We formulate the relation and frame-argument extraction as a single dependency parsing task, which extracts *all* relations in a sentence in one processing step. (4) We perform an extensive set of multi-task learning experiments with related corpora, showing that MuLMS is a useful auxiliary task for two other materials science NLP datasets.

2 Related Work

Several **materials science NLP datasets** have recently been released, e.g., targeting NE recognition (Yamaguchi et al., 2020; O’Gorman et al., 2021). The Materials Science Procedural Text (MSPT) corpus (Mysore et al., 2019) consists of paragraphs describing synthesis procedures annotated with graph structures capturing relations and typed arguments. SOFC-Exp (Friedrich et al., 2020) marks similar graph structures describing experiments.

In this paper, we compare two state-of-the-art approaches to **Named Entity Recognition** (NER). Huang et al. (2015) use a CRF layer (Lafferty et al., 2001) on top of a neural language model (in their case a BiLSTM) for sequence-tagging related tasks. Yu et al. (2020) treat NER as a graph-based dependency parsing task by representing NEs as spans between the first and last token of an entity. In the materials science domain, Friedrich et al. (2020) test a variety of embedding combinations in a CRF-based tagger. O’Gorman et al. (2021) compare different pre-trained transformers for token classification. Both studies find SciBERT (Beltagy et al., 2019), a BERT-style model pre-trained on scientific documents, to be very effective.

Relation and Event Extraction. Friedrich et al.

(2020) treat their slot filling task in the SOFC sub-domain as a sequence tagging task, assuming that each sentence represents at most one experiment. To predict a possible relation between two entities, Swarup et al. (2020) retrieve a set of sentences similar to the input sentence, and learn to copy relations from these sentences. Mysore et al. (2017) experiment with unsupervised methods for extracting action graphs for synthesis procedures.

An exhaustive overview of the literature on biomedical relation extraction is out of the scope of this paper. Recent works have used graph-neural networks (Huang et al., 2020), or convolutional neural networks (Ramponi et al., 2020). Sarrouti et al. (2022) compare various pre-trained transformer models. Semantic parsing of frame structures (Fillmore and Baker, 2001) has been addressed using graph-convolutional networks (Marcheggiani and Titov, 2020), BiLSTMs (He et al., 2018), and recently by generating structured output using encoder-decoder models (Hsu et al., 2022; Lu et al., 2021). Tackling semantic dependency parsing with a biaffine classifier architecture was first proposed by Dozat and Manning (2018).

3 MuLMS Corpus

In this section, we present our new corpus.

3.1 Source of Texts and Preprocessing

We select 50 scientific articles licensed under CC BY from seven popular sub-areas of materials science: electrolysis, graphene, polymer electrolyte fuel cell (PEMFC), solid oxide fuel cell (SOFC), polymers, semiconductors, and steel. The four SOFC papers were selected from the SOFC-Exp corpus (Friedrich et al., 2020). 11 papers were selected from the OA-STM corpus² and classified into the above subject areas by a domain ex-

¹<https://github.com/boschresearch/mulms-wiesp2023>

²<https://github.com/elsevierlabs/OA-STM-Corpus>

pert. The majority of the papers were found via PubMed³ and DOAJ⁴ using queries prepared by a domain expert. For the OA-STM data, we use the sentence segmentation provided with the corpus, which has been created using GENIA tools (Tsuruoka and Tsujii, 2005). For the remaining texts, we rely on the sentence segmentation provided by INCEpTION (Klie et al., 2018) with some manual fixes. As shown in Table 1, documents are rather long with a tendency to long sentences (but with high variation due to, i.a., short headings).

3.2 Annotation Scheme

We annotate various layers: NEs, relations, and frame structures representing measurements.

3.2.1 Named Entities

We annotate the following materials-science specific NE mentions and assign the following NE types to these mentions:

MAT: mentions of materials described by their chemical formula (WO_3) or its chemical name (*indium tin oxide*).

FORM: mentions of the form or morphology of the material, e.g., *thin film, gas, liquid, cubic*.

INSTRUMENT: mentions of devices used to perform a materials-science-related measurement, e.g., *Olympus BX52 microscope*.

DEVICE: mentions of devices (target products) whose construction or improvement is the aim of the research (e.g., *photodetector, transistor, supercapacitor*). DEVICE is not used for instrumentation devices that are only used as a tool.

NUM: mentions of numbers such as *0.46*.

UNIT: mentions of units such as *nm* or *V*.

RANGE: mentions of numeric expressions indicating ranges, e.g., *0.46±.11*

VALUE: nested type capturing expressions of values usually composed of a NUM, RANGE and a UNIT, e.g., *~5 x 3mm²*

CITE: citations, e.g., *Setman et al.* or *[13]*.

PROPERTY: expressions referring to properties of materials or conditions in experiments, e.g., *stress rate* or *electron conductivity*.

TECHNIQUE: mentions of experimental techniques used in the characterization steps, e.g., *Scanning electron microscopy (SEM)*.

SAMPLE: mention of the material or a component made of materials studied in a measurement, either referred to by a particular name or its com-

#Documents	50
#Documents train / dev / test	36 / 7 / 7
#Sentences	10186
#Sentences/Document	203.7±73.2
#Tokens/Sentence	28.7±17.9

Table 1: Basic **corpus statistics** for MuLMS.

position, its batch name (*Aq-825*) or by referring to the whole component (*MEA-Pt/C*) or to part of the material’s structure (*ionomer patches*). In simulation papers, the SAMPLE may also be the computational model under study (*RBF-ANN*).

3.2.2 Relations and Measurement Frame

We treat measurement annotation in a frame-like (Fillmore and Baker, 2001) fashion, using the span type MEASUREMENT to mark the triggers (e.g., *was measured, is plotted*) that introduce the *Measurement* frame to the discourse. About 88% of the triggers are verbs. The remaining 12% occur in figure captions without verb phrases and are annotated either on nouns (*Comparison*) or, in absence of more suitable phrases, on figure labels such as *Figure 17*. The trigger annotations of these sentences serve as the root of the tree/graph annotations as illustrated in Figure 1.

There are also cases in which the *Measurement* frame is evoked, but there are no technical details or results that we can extract about the measurement. We mark the triggers of these sentences with the tag QUAL_MEAS (qualitative mention of a measurement). An example of such a sentence is “*We compare a critical volume to be detached from the different nanostructures.*”

Measurement-related Relations. We annotate several relations that start at a MEASUREMENT tag and that end at the annotations of the corresponding slot fillers within the sentence. Consider the following sentence: “*To characterize the ORR activity of the catalyst, linear scan voltammetry (LSV) was tested from 0 to 1.2 V on an RDE with a scan rate of 50 mV/s in O₂-saturated HClO₄ solution.*”

measuresProperty: indicates the PROPERTY (e.g., *ORR activity*) that is measured.

conditionSampleFeatures: indicates the SAMPLE or MATERIAL whose property is measured. In the above example, the sample is the *catalyst*.

usesTechnique: relates to the TECHNIQUE (e.g., *linear scan voltammetry*) used in a measurement.

conditionInstrument: refers to the INSTRUMENT used to make a measurement, e.g., *RDE/rotating*

³<https://pubmed.ncbi.nlm.nih.gov>

⁴<https://doaj.org>

Relation	Example
<i>hasForm</i>	silicon _{MAT} <i>hasForm</i> → hexagonal _{FORM}
<i>usedIn</i>	SiC _{MAT} <i>usedIn</i> → MOSFET _{DEVICE}
<i>usedAs</i>	PtNi3M _{MAT} <i>usedAs</i> → catalysts
<i>dopedBy</i>	chlorinated _{MATERIAL} <i>dopedBy</i> → SiC _{MAT}

Table 2: Measurement-independent **relations** annotated in MuLMS. MAT is short for MATERIAL.

Label	Count	%	Label	Count	%
MAT	15596	33.6	CITE	1709	3.7
NUM	6081	13.1	SAMPLE	1461	3.2
VALUE	4852	10.5	TECHNIQUE	1036	2.2
UNIT	4330	9.3	DEV	808	1.7
PROPERTY	3925	8.5	RANGE	736	1.6
FORM	3568	7.7	INSTRUMENT	378	0.8
MEASUREMENT	2171	4.7	<i>total</i>	46,351	-

Table 3: Corpus counts for **named entity** annotations.

disk electrode.

conditionProperty: a property that is a condition in the experiment, e.g., *scan rate* (which in turn has the *propertyValue* of *50 mv/s*).

propertyValue: connects the mention of a PROPERTY and that of its corresponding VALUE. This relation may also occur if a mention of a PROPERTY occurs independently of a measurement.

conditionEnvironment: identifies the MATERIALS (e.g., *O2* and *HClO4*) and VALUES (e.g., an operating temperature of *30°C*) that provide the environment of the measurement.

takenFrom: connects the MEASUREMENT with the bibliographic reference CITE from which the setup has been inspired or taken over.

In most cases, a *conditionProperty* or a *measuresProperty* connects the MEASUREMENT annotation to a PROPERTY node, at which a *propertyValue* relation starts that ends at the respective VALUE. However, in some cases, the condition or measured property is not mentioned explicitly. In this case, we link the VALUE directly to the MEASUREMENT node via a *conditionPropertyValue* or a *measuresPropertyValue* link. For consistency reasons, we also add these links in cases that mention the property explicitly, turning the trees into graph structures. Out of the added *conditionPropertyValue* links, 967 are for such explicit cases, while the other 206 describe implicit cases. In the case of *measuresPropertyValue*, 722 links are for explicit cases and 36 for implicit cases.

Further Relations. In the following, we explain relations that can appear independently of measure-

Label	Count	%	Label	Count	%
<i>hasForm</i>	2910	17.3	<i>meas.Prop.Val.</i>	751	4.5
<i>measuresProperty</i>	2080	12.4	<i>usedTogether</i>	672	4.0
<i>usedAs</i>	1839	11.0	<i>conditionEnv.</i>	549	3.3
<i>propertyValue</i>	1794	10.7	<i>usedIn</i>	434	2.6
<i>conditionProperty</i>	1648	9.8	<i>conditionInstr.</i>	357	2.2
<i>conditionSample</i>	1434	8.5	<i>takenFrom</i>	118	0.7
<i>cond.Prop.Value</i>	1158	6.9	<i>dopedBy</i>	65	0.4
<i>usesTechnique</i>	985	5.9	<i>total</i>	16,794	-

Table 4: Corpus counts for **measurement relations**.

	MEAS	QUAL_MEAS	OTHER
MEAS	48	6	6
QUAL_MEAS	12	37	11
OTHER	10	17	92

Table 5: Inter-annotator agreement for **identifying measurement sentences**: confusion matrix.

ments. Examples are shown in Table 2.

hasForm: connects mention of MATERIAL and the corresponding FORM annotation.

usedIn: connects MATERIAL and the DEVICE it is used in. In Table 2, *MOSFET* stands for *Metal Oxide Semiconductor Field-Effect Transistors*.

usedAs: links a specific MATERIAL mention with a more generic one such as *catalyst*, a material class defined by its function.

dopedBy: indicates dopants (e.g., *chlorine*), i.e., impurities added to a main material (e.g., *SiC*).

usedTogether: connects two MATERIALS if they are used together in an experiment, i.e., if the materials are part of an assembly or a mixture.

3.3 Corpus Statistics

We now analyze our corpus and provide detailed corpus statistics. In total, there are 46,351 NE annotations. Table 3 shows the counts by NE label. There are roughly 1.5 MAT annotations per sentence as these are nested and occurrences of composite materials often result in many combined MAT tags. Table 4 reports the counts of annotated relations (16,794 in total), with *hasForm* as the most frequent relation with 2910 instances and *dopedBy* the least frequent with only 65 instances.

Out of all 10186 sentences, 2111 (20.7%) describe a measurement (i.e., they contain at least one MEASUREMENT annotation). On average, each document contains 43.4 MEASUREMENT annotations. In addition, there are 1476 sentences (14.5%) marked as containing a QUAL_MEAS, with 40 sentences of these also containing a MEASUREMENT annotation.

Label	P	R	F1	Label	P	R	F1
MAT	96.7	91.2	93.9	CITE	97.4	97.4	97.4
NUM	98.9	100.0	99.4	SAMPLE	3.1	12.5	4.7
VALUE	100	100	100.0	TECHN.	77.5	59.6	67.4
UNIT	97.9	100.0	98.9	DEVICE	96.0	82.8	88.9
PROPERTY	42.6	37.7	40.0	RANGE	100.0	100.0	100.0
FORM	95.7	86.3	90.8	INSTR.	80.0	76.9	78.4
MEAS.	44.6	51.0	47.6	<i>average</i>	<i>79.3</i>	<i>76.6</i>	<i>77.9</i>

Table 6: Inter-annotator agreement: **named entities**.

3.4 Inter-Annotator Agreement (IAA)

Our entire dataset has been annotated by a graduate student of materials science, who was also involved in the design of the annotation scheme. We perform two agreement studies, comparing to the annotations of a second annotator with a PhD degree in environment engineering and several years of experience in materials science research.

Agreement on identifying Measurement sentences. In this agreement study, we estimate the degree of agreement whether a sentence expresses a MEASUREMENT, a QUAL_MEAS, or whether it does not express a measurement at all. We sample 60 sentences marked with MEASUREMENT, 60 sentences marked with QUAL_MEAS, and 120 sentences not marked as either by the first annotator. Table 5 shows the confusion matrix for the 239 sentences for which both annotators provided a label. One automatically selected sentence was not labeled by one of the annotators due to incomprehensibility. In terms of Cohen’s κ (Cohen, 1960), agreement amounts to 59.2, indicating moderate to substantial agreement (Landis and Koch, 1977). When collapsing MEASUREMENT and QUAL_MEAS, κ is 63.4 (substantial).

Agreement on named entities. We next compute agreement for NE and relation annotations. IAA on “easy” types such as MAT, NUM, UNIT, VALUE and RANGE has been shown to be very high in prior work (Friedrich et al., 2020). Hence, as our resources are limited, we provide annotations of these types for correction to the second annotator. We sample 134 sentences such that each entity type occurs at least 25 times in the annotations of the first annotator and have the second annotator correct or add entity annotations. We then compare the annotated sets of NE mentions using precision and recall (for a justification of this choice of agreement metrics, see Appendix C). Results using relaxed matching (containment) are shown in Table 6 (detailed counts in Appendix C).

Label	P	R	F1	κ	matches
<i>propertyValue</i>	81.2	81.2	81.2	0.88	26
<i>condSampleFeat.</i>	43.1	40.0	41.5	0.23	22
<i>usedIn</i>	40.0	52.2	45.3	0.55	12
<i>usesTechnique</i>	72.9	67.3	70.0	0.86	35
<i>hasForm</i>	54.7	71.4	61.9	0.64	35
<i>takenFrom</i>	33.3	63.6	43.7	0.57	7
<i>measuresProp.</i>	80.5	72.9	76.5	0.80	70
<i>dopedBy</i>	35.3	50.0	41.4	0.25	6
<i>conditionProp.</i>	27.7	59.1	37.7	0.31	13
<i>conditionInstr.</i>	24.0	60.0	34.3	0.44	6
<i>conditionEnv.</i>	0.0	0.0	0.0	0.00	0
<i>usedAs</i>	23.5	85.7	36.9	0.37	12
<i>usedTogether</i>	13.8	28.6	18.6	0.20	4

Table 7: Inter-annotator agreement: **relations**.

For most types, scores are in the expected range of difficult semantic annotation tasks. Agreement on identifying Measurement sentences is good; the decision of where exactly to place the MEASUREMENT annotation differs between annotators.

Agreement on relations. We sample 178 sentences in which each relation occurs at least 25 times according to the first annotator. We keep NE annotations and ask the second annotator to add relations. Table 7 shows the results in terms of precision, recall, and κ per relation type. The latter has been computed by treating all pairs of NE annotations as potential relations, using *NO_REL* if no relation has been annotated. Overall κ on relations is 0.61 (substantial). For each relation label, we can map all other relation types to *OTHER* and compute agreement for the binary decision whether the label is present or not (analysis suggested by Krippendorff (1989)). κ aims to quantify the degree of agreement *above* chance. Interpreting our κ scores according to the scale of Landis and Koch (1977), we reach at **fair** agreement for *conditionPropertyValue*, *usedTogether*, *conditionSampleFeatures*, *dopedBy*, and *usedAs*. We reach **moderate** agreement for *usedIn*, *takenFrom*, and *conditionInstrument*. For the practically important relations *propertyValue*, *measuresProperty*, and *usesTechnique*, we even reach **almost perfect** agreement.

For the non-easily identifiable types, post-hoc discussion with the second annotator (who did not receive an extensive training on the task) concluded it was not always clear to them when using related labels (e.g., *conditionProperty* and *conditionEnvironment*). Yet, these labels can be learned with good or acceptable accuracy (see Appendix E), indicating that the primary annotator has used the labels consistently.

4 Task Definitions and Modeling

In this section, we define several NLP tasks for MuLMS and describe our computational models.

4.1 Pre-trained Models

We use BERT (Devlin et al., 2019) as the underlying text encoder for all of our models. We also use variants of BERT, namely SciBERT (Beltagy et al., 2019), which has been pre-trained on articles in the scientific domain, and MatSciBERT (Gupta et al., 2022), a version of SciBERT further pre-trained on materials science articles. We use the uncased, 768-dimensional variant of each model, which we fine-tune.

4.2 Detecting Measurements

We model the task of classifying whether a sentence contains a MEASUREMENT or a QUAL_MEAS annotation as a ternary sentence classification task, i.e., it is also possible that a sentence does not refer to any measurement. As we are primarily interested in detecting MEASUREMENT, we map the few multi-label cases carrying both positive labels to MEASUREMENT. We use a linear layer plus softmax with the CLS token embedding as input. For training, we downsample the amount of non-measurement sentences.

4.3 Named Entity Recognition (NER)

We compare two state-of-the-art models for NER, (a) a sequence tagger and (b) a dependency parser. For the **sequence tagger**, we encode the NE labels using the nested BILOU scheme (Alex et al., 2007), which leverages a label set of combined types constructed from the training set for nested NEs. As there are only very few cases (about 0.65% of all NE annotations) where a token receives more than three stacked NE labels, in order to avoid sparsity issues, we consider only the “bottom” three layers of stacked entities. We feed the contextualized embeddings of the last transformer layer of the respective first wordpiece token of each “real” token into a linear layer and then use a CRF (Lafferty et al., 2001) to optimize predictions for the entire sequence.

Modeling NER as a **dependency parsing** task (Yu et al., 2020) can easily account for nested NEs. The main idea is to predict edges reaching from the end token of an NE to its start token as depicted in Figure 2. We adapt the STEPS parsing pipeline (Grünwald et al., 2021a) to the

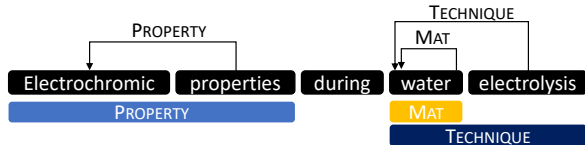


Figure 2: NER as dependency parsing (Yu et al., 2020).

task. There are three combinations of tags in our dataset that occasionally cover the exact same span and that occur more than 20 times: VALUE+NUM, VALUE+RANGE and MAT+FORM. We hence introduce the above combined labels. For any other infrequent conflicting labels, we do not add extra tags, i.e., the model can never catch these cases. We decide on this slight restriction of the model capabilities in order to avoid sparsity issues. In the evaluation, we do not filter for these cases but of course use all nested NEs as annotated.

4.4 Relation Extraction

Given an input sentence along with all named entities within it, as well as their types (either gold or predicted depending on the experimental setting), we predict which (if any) relation is present between them. We treat all relations in a single model and predict all relations of a sentence simultaneously by modeling relation extraction as a graph parsing task. Following Toshniwal et al. (2020), we first create an embedding e_i for the i^{th} NE in the sentence by concatenating the token embeddings of its first and last token ($e_{i,\text{START}}, e_{i,\text{END}}$). We also concatenate a learned embedding for the NE’s label ($e_{i,\text{LABEL}}$): $e_i = e_{i,\text{START}} \oplus e_{i,\text{END}} \oplus e_{i,\text{LABEL}}$

Considering NEs as nodes in a graph, we use a biaffine classifier architecture (Dozat and Manning, 2017) using the implementation of Grünwald et al. (2021a,b) to predict the relation between each pair. The non-existence of a relation is encoded as simply another label (\emptyset). For details on the parser architecture, see Appendix B.

5 Experiments

We now detail our experimental results.

Experimental Settings. We split our corpus into train, dev, and test sets on a per-document basis. Within the train set, we provide five distinct tune splits (*train1* to *train5*). For all experiments and for hyperparameter tuning, we always train five models. Similar to cross-validation, we train on four folds and use the fifth “training fold” for model selection (cf. van der Goot (2021) for details). Hy-

LM (sampling)	MEASUREMENT			QUAL_MEAS		
	P	R	F1	P	R	F1
<i>Random baseline</i>	24.7	19.7	21.9 \pm 2.4	15.1	14.4	14.7 \pm 1.5
BERT(0.7)	74.1	71.4	72.5 \pm 2.1	49.9	51.6	50.6 \pm 0.7
SciBERT(0.7)	71.1	79.5	75.0 \pm 0.7	52.7	52.8	52.7 \pm 1.4
MatSciBERT(0.85)	70.6	80.1	74.9 \pm 1.3	52.8	56.9	54.7 \pm 1.0
<i>human agreement*</i>	68.6	80.9	73.8	61.7	61.7	61.7

Table 8: Ternary sentence classification results for **identifying measurement sentences** on test set. ‘‘Sampling’’ indicates amount of OTHER sentences used for training. **estimated on subset of data.*

Model	LM	Micro F1	Macro F1
Dependency Parser	BERT	73.0 \pm 0.7	61.7 \pm 2.3
	SciBERT	76.5 \pm 0.3	65.8 \pm 0.9
	MatSciBERT	77.3 \pm 0.3	67.6 \pm 1.4
CRF Tagger	BERT	75.3 \pm 0.6	63.2 \pm 0.7
	SciBERT	78.7 \pm 0.4	69.3 \pm 1.0
	MatSciBERT	79.6 \pm 0.4	70.7 \pm 0.7

Table 9: **Named entity recognition** results on test set.

perparameters are chosen based on the best dev results, and we finally report results for the test set. The splits are the same across all tasks. Because the training data varies across the five runs for which we report results, standard deviations are usually larger than when using the same training data. For hyperparameter settings, see Appendix A.

5.1 Identifying Measurement Sentences

Table 8 reports the results for identifying sentences that contain a MEASUREMENT or a QUAL_MEAS annotation. In each experiment, we tune the down-sampling rate for the majority class OTHER and the learning rate (using grid search from 1e-4 to 1e-7). The *random baseline* assigns labels according to the percentage of instances in the (full) training set carrying a particular label. The average overall accuracy of the MatSciBERT classifier is 78.2%. SciBERT and MatSciBERT perform similarly, with MatSciBERT having a small edge. Identification of MEASUREMENT is comparable to our estimate of human agreement. For identifying QUAL_MEAS, there is headroom.

5.2 Named Entity Recognition Results

Table 9 shows the results for named entity recognition. Again, MatSciBERT performs best with Micro F1 scores approaching 80, which indicates that NE mentions are consistently annotated in MuLMS. The CRF-based tagger outperforms the dependency-parser-based NER model by a consid-

LM	dev		test	
	Micro F1	Macro F1	Micro F1	Macro F1
<i>Maj. basel.</i>	38.3	29.4	37.2	27.4
BERT	69.5 \pm 0.5	63.4 \pm 1.0	63.5 \pm 0.6	57.7 \pm 1.1
SciBERT	72.5 \pm 0.8	65.7 \pm 0.6	67.5 \pm 0.9	62.0 \pm 2.2
MatSciBERT	73.2 \pm 1.0	66.5 \pm 1.1	67.6 \pm 1.0	62.0 \pm 1.0

Table 10: **Relation extraction** results: gold entities.

erable margin. For detailed per-label statistics, see Appendix E. Precision and recall are approximately balanced for all labels. An exception is SAMPLE, which is both infrequent in the dataset and hard to identify for humans. Both models suffer from low recall for this tag.

5.3 Relation Extraction Results

Table 10 shows the results for relation extraction on gold entities. A predicted relation is counted as correct if and only if there is a relation with the same start span, end span, and relation label in the set of gold relations for the sentence. The majority baseline assigns to each pair of entities the relation that is most common in the training set for the respective entity types of the governing and dependent spans (see Appendix E).

The results demonstrate that a biaffine dependency parsing approach achieves robust performance overall and outperforms the baseline by a substantial margin. The two models trained on scientific text outperform BERT. Their results are similar, with MatSciBERT having a slight edge.

Analysis of per-label scores (see Appendix E) for MatSciBERT) shows that the highest scores are achieved for *conditionInstrument* (92.2 F1), *usesTechnique* (91.0 F1), and *takenFrom* (84.7). This is somewhat surprising especially for *conditionInstrument* and *takenFrom*, as these are among the rarest relation types in the corpus (see Table 4). However, our majority baseline achieves high accuracies on these relation types as well (>90 F1 for *conditionInstrument* and *usesTechnique*), i.e., they are easily inferable from entity types. The worst performance is observed on the relation types *usedTogether* (4.0 F1), *dopedBy* (22.7 F1), and *usedIn* (37.9 F1). These relations occur relatively rarely and also cannot be inferred from entity types.

Relation extraction on predicted entities. Finally, we also run our relation extraction module on predicted named entities using the respective best-performing models (both based on MatSciBERT). Models are evaluated as above, with the additional

Dataset	LR	Micro F1	Macro F1
MuLMS	1e-4/7e-3	79.6 \pm 0.4	70.7 \pm 0.7
+ SOFC-Exp		79.5 \pm 0.4	70.4 \pm 1.3
+ MSPT		79.2 \pm 0.6	69.9 \pm 0.8
SOFC-Exp	3e-4/7e-3	83.4 \pm 0.9	81.0 \pm 1.0
+ MuLMS		81.9 \pm 1.7	79.6 \pm 1.7
MSPT	5e-5/9e-3	81.6 \pm 0.4	57.8 \pm 0.6
+ MuLMS		80.4 \pm 0.4	56.4 \pm 0.8

Table 11: **NER MTL** results: MatSciBERT tagger.

requirement that the boundaries of start/end spans of a predicted relation must also exactly match those of the respective gold spans. Prediction accuracy drops substantially: to micro-F1 scores of 42.5 and 36.5 on dev and test, respectively, corresponding to macro-F1 scores of 37.9 and 32.8. The reason for this is error propagation as relations can only be retrieved if the entities are predicted correctly, and as incorrectly labeled entities can mislead the relation classifier.

5.4 Multi-Task Learning Across Datasets

To find out whether information extraction accuracy can be increased by employing multi-task learning (MTL), we perform a series of experiments in which we combine MuLMS training data with NE and relation data from other materials science datasets, namely the SOFC-Exp and MSPT corpora (see Sec. 2).⁵ In all experiments, we use a shared MatSciBERT and one classification head for each task (dataset). When reporting results on MuLMS, we use the same setup as before, but add the complete training sets of SOFC-Exp or MSPT during training. When reporting results on SOFC and MSPT, we train on all of their training data and the complete training data of MuLMS and perform early stopping on dev. For these experiments, reported scores are averages over 5 runs with different random seeds.

For **NER** (Table 11), we do not observe overall improvements. We hypothesize that this is because in SOFC-Exp, NE types are much coarser-grained, and in MSPT, NE annotations are focused on synthesis procedure paragraphs only. Nevertheless, as can be seen by the per-label scores in Appendix E, average scores on MuLMS are mainly hurt by decreases on SAMPLE, while scores for RANGE increase considerably by up to 3.9%.

⁵We removed one document from the test split of the SOFC-Exp corpus that is also part of the train set of MuLMS.

Dataset	dev		test	
	Micro F1	Macro F1	Micro F1	Macro F1
MuLMS	73.2 \pm 1.0	66.5 \pm 1.1	67.6 \pm 1.0	62.0 \pm 1.0
+ SOFC	72.5 \pm 0.9	65.8 \pm 0.7	68.1 \pm 0.7	61.1 \pm 0.7
+ MSPT	73.9 \pm 0.4	67.4 \pm 0.4	68.7 \pm 0.7	63.7 \pm 0.7
SOFC-Exp	71.3 \pm 0.6	62.8 \pm 2.1	66.9 \pm 1.6	59.8 \pm 1.5
+ MuLMS	72.3 \pm 0.5	64.3 \pm 3.7	68.7 \pm 1.5	60.9 \pm 3.6
+ MSPT	72.3 \pm 0.9	63.1 \pm 2.6	67.6 \pm 2.0	60.8 \pm 4.3
MSPT	84.2 \pm 0.6	82.5 \pm 0.7	84.6 \pm 0.8	83.0 \pm 0.8
+ MuLMS	85.3 \pm 0.2	83.4 \pm 0.8	85.6 \pm 0.4	84.1 \pm 0.6
+ SOFC	83.7 \pm 1.4	81.8 \pm 1.4	84.7 \pm 1.1	83.2 \pm 1.4

Table 12: Relation extraction multi-tasking results using MatSciBERT-based parser.

Results for MTL for **relations** are shown in Table 12. We observe that adding MuLMS to the training data of both SOFC-Exp and MSPT results in improvements. Incorporating SOFC-Exp instances in the training does not meaningfully increase prediction accuracy on MuLMS, whereas incorporating instances from MSPT leads to modest improvements. Intuitively, this makes sense: relations in SOFC-Exp focus on a specific type of experiment, while MuLMS covers a broader range of measurements. Similarly, some MuLMS relations bear resemblance to MSPT relations (e.g., those dealing with instruments or apparatus), which explains why training jointly is beneficial.

6 Conclusion and Outlook

In this resource paper, we have presented a new large-scale dataset of 50 scientific articles in the domain of materials science *exhaustively* annotated with named entity mentions, relations, and measurement-related frames. Our inter-annotator agreement study shows good agreement for most decisions. Our experiments with state-of-the-art neural models highlight that most distinctions can be learned with good accuracy, and that synergies can be achieved by training jointly with existing more specific materials-science NLP datasets.

Future work is needed to improve on end-to-end or joint models of NER and relation extraction as our experiments showed that a pipeline-based setting suffers from error propagation. A potential next step is to adapt sequence-to-sequence models to the structure induction tasks of MuLMS, following ideas of (Hsu et al., 2022; Lu et al., 2021). Finally, employing data augmentation techniques in particular for the less frequent relation types is a viable path for future work.

Limitations

As discussed in Sec. 3.4, we expect our inter-annotator agreement scores to underestimate the reproducibility of the task. It is, unfortunately, not trivial to find annotators with the required background knowledge. Hence, scores reflect agreement after only an initial very brief training phase, but nevertheless (in our opinion) give useful insights on the relative difficulty of the labeling decisions.

In our relation extraction experiments, we use label embeddings based on either gold or predicted entity labels (depending on the experimental setup) as an input to our system. Providing gold entity label information in particular constitutes a setting that is considerably easier for a relation classifier than providing no label information. Using predicted entity mention and labels showed to suffer from error propagation. In future work, it may be interesting to evaluate the performance of a relation extraction system that is not given label information, or that predicts entity labels jointly with relations.

Ethical Considerations

The annotators participating in our project were completely aware of the goal of the annotations and even helped designing the annotation scheme. They gave explicit consent to the publication of their annotations. The main annotator was paid considerably above our country's minimum wage.

References

- Beatrice Alex, Barry Haddow, and Claire Grover. 2007. [Recognising nested named entities in biomedical text](#). In *Biological, translational, and clinical language processing*, pages 65–72, Prague, Czech Republic. Association for Computational Linguistics.
- Iz Beltagy, Arman Cohan, Guy Feigenblat, Dayne Freitag, Tirthankar Ghosal, Keith Hall, Drahomira Herrmannova, Petr Knoth, Kyle Lo, Philipp Mayr, Robert M. Patton, Michal Shmueli-Scheuer, Anita de Waard, Kuansan Wang, and Lucy Lu Wang, editors. 2021. [Proceedings of the Second Workshop on Scholarly Document Processing](#). Association for Computational Linguistics, Online.
- Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. [SciBERT: A pretrained language model for scientific text](#). 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 3615–3620, Hong Kong, China. Association for Computational Linguistics.
- Alex Brandsen, Suzan Verberne, Milco Wansleeben, and Karsten Lambers. 2020. [Creating a dataset for named entity recognition in the archaeology domain](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 4573–4577, Marseille, France. European Language Resources Association.
- Muthu Kumar Chandrasekaran, Anita de Waard, Guy Feigenblat, Dayne Freitag, Tirthankar Ghosal, Edward Hovy, Petr Knoth, David Konopnicki, Philipp Mayr, Robert M. Patton, and Michal Shmueli-Scheuer, editors. 2020. [Proceedings of the First Workshop on Scholarly Document Processing](#). Association for Computational Linguistics, Online.
- Arman Cohan, Guy Feigenblat, Dayne Freitag, Tirthankar Ghosal, Drahomira Herrmannova, Petr Knoth, Kyle Lo, Philipp Mayr, Michal Shmueli-Scheuer, Anita de Waard, and Lucy Lu Wang, editors. 2022. [Proceedings of the Third Workshop on Scholarly Document Processing](#). Association for Computational Linguistics, Gyeongju, Republic of Korea.
- Jacob Cohen. 1960. [A coefficient of agreement for nominal scales](#). *Educational and Psychological Measurement*, 20(1):37–46.
- Kevin B. Cohen, Dina Demner-Fushman, Sophia Ananiadou, Bonnie Webber, Jun'ichi Tsujii, and John Pestian, editors. 2012. [BioNLP: Proceedings of the 2012 Workshop on Biomedical Natural Language Processing](#). Association for Computational Linguistics, Montréal, Canada.
- Nigel Collier, Patrick Ruch, and Adeline Nazarenko, editors. 2004. [Proceedings of the International Joint Workshop on Natural Language Processing in Biomedicine and its Applications, NLPBA/BioNLP 2004, Geneva, Switzerland, August 28-29, 2004](#).
- Dina Demner-Fushman, Kevin Bretonnel Cohen, Sophia Ananiadou, and Junichi Tsujii, editors. 2022. [Proceedings of the 21st Workshop on Biomedical Language Processing](#). Association for Computational Linguistics, Dublin, Ireland.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Timothy Dozat and Christopher D. Manning. 2017. [Deep biaffine attention for neural dependency parsing](#). In *International Conference on Learning Representations*.

- Timothy Dozat and Christopher D. Manning. 2018. [Simpler but more accurate semantic dependency parsing](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 484–490, Melbourne, Australia. Association for Computational Linguistics.
- Charles J Fillmore and Collin F Baker. 2001. Frame semantics for text understanding. In *Proceedings of WordNet and Other Lexical Resources Workshop, NAACL*, volume 6.
- Annemarie Friedrich, Heike Adel, Federico Tomazic, Johannes Hingerl, Renou Benteau, Anika Maruszczyk, and Lukas Lange. 2020. [The SOFC-exp corpus and neural approaches to information extraction in the materials science domain](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1255–1268, Online. Association for Computational Linguistics.
- Stefan Grünewald, Annemarie Friedrich, and Jonas Kuhn. 2021a. [Applying occam’s razor to transformer-based dependency parsing: What works, what doesn’t, and what is really necessary](#). In *Proceedings of the 17th International Conference on Parsing Technologies and the IWPT 2021 Shared Task on Parsing into Enhanced Universal Dependencies (IWPT 2021)*, pages 131–144, Online. Association for Computational Linguistics.
- Stefan Grünewald, Frederik Tobias Oertel, and Annemarie Friedrich. 2021b. [RobertNLP at the IWPT 2021 shared task: Simple enhanced UD parsing for 17 languages](#). In *Proceedings of the 17th International Conference on Parsing Technologies and the IWPT 2021 Shared Task on Parsing into Enhanced Universal Dependencies (IWPT 2021)*, pages 196–203, Online. Association for Computational Linguistics.
- Tanishq Gupta, Mohd Zaki, N. M. Anoop Krishnan, and Mausam. 2022. [MatSciBERT: A materials domain language model for text mining and information extraction](#). *npj Computational Materials*, 8(1):102.
- Luheng He, Kenton Lee, Omer Levy, and Luke Zettlemoyer. 2018. [Jointly predicting predicates and arguments in neural semantic role labeling](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 364–369, Melbourne, Australia. Association for Computational Linguistics.
- George Hripcsak and Adam S Rothschild. 2005. [Agreement, the f-measure, and reliability in information retrieval](#). *Journal of the American medical informatics association*, 12(3):296–298.
- I-Hung Hsu, Kuan-Hao Huang, Elizabeth Boschee, Scott Miller, Prem Natarajan, Kai-Wei Chang, and Nanyun Peng. 2022. [DEGREE: A data-efficient generation-based event extraction model](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1890–1908, Seattle, United States. Association for Computational Linguistics.
- Kung-Hsiang Huang, Mu Yang, and Nanyun Peng. 2020. [Biomedical event extraction with hierarchical knowledge graphs](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1277–1285, Online. Association for Computational Linguistics.
- Zhiheng Huang, Wei Xu, and Kai Yu. 2015. [Bidirectional LSTM-CRF models for sequence tagging](#). In *CorRR*, volume abs/1508.01991.
- Jan-Christoph Klie, Michael Bugert, Beto Boullosa, Richard Eckart de Castilho, and Iryna Gurevych. 2018. [The INCEpTION platform: Machine-assisted and knowledge-oriented interactive annotation](#). In *Proceedings of the 27th International Conference on Computational Linguistics: System Demonstrations*, pages 5–9, Santa Fe, New Mexico. Association for Computational Linguistics.
- Dan Kondratyuk and Milan Straka. 2019. [75 languages, 1 model: Parsing Universal Dependencies universally](#). 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 2779–2795, Hong Kong, China. Association for Computational Linguistics.
- Klaus Krippendorff. 1989. Content analysis.
- John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. 2001. [Conditional random fields: Probabilistic models for segmenting and labeling sequence data](#). In *Proceedings of the Eighteenth International Conference on Machine Learning (ICML 2001)*, Williams College, Williamstown, MA, USA, June 28 - July 1, 2001, pages 282–289. Morgan Kaufmann.
- J Richard Landis and Gary G Koch. 1977. An application of hierarchical kappa-type statistics in the assessment of majority agreement among multiple observers. *Biometrics*, pages 363–374.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#).
- Yaojie Lu, Hongyu Lin, Jin Xu, Xianpei Han, Jialong Tang, Annan Li, Le Sun, Meng Liao, and Shaoyi Chen. 2021. [Text2Event: Controllable sequence-to-structure generation for end-to-end event extraction](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2795–2806, Online. Association for Computational Linguistics.
- Diego Marcheggiani and Ivan Titov. 2020. [Graph convolutions over constituent trees for syntax-aware semantic role labeling](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language*

- Processing (EMNLP)*, pages 3915–3928, Online. Association for Computational Linguistics.
- Sheshera Mysore, Zachary Jensen, Edward Kim, Kevin Huang, Haw-Shiuan Chang, Emma Strubell, Jeffrey Flanigan, Andrew McCallum, and Elsa Olivetti. 2019. [The materials science procedural text corpus: Annotating materials synthesis procedures with shallow semantic structures](#). In *Proceedings of the 13th Linguistic Annotation Workshop*, pages 56–64, Florence, Italy. Association for Computational Linguistics.
- Sheshera Mysore, Edward Kim, Emma Strubell, Ao Liu, Haw-Shiuan Chang, Srikrishna Kompella, Kevin Huang, Andrew McCallum, and Elsa Olivetti. 2017. [Automatically extracting action graphs from materials science synthesis procedures](#). *NIPS Workshop on Machine Learning for Molecules and Materials*.
- Tim O’Gorman, Zach Jensen, Sheshera Mysore, Kevin Huang, Rubayat Mahbub, Elsa Olivetti, and Andrew McCallum. 2021. [MS-mentions: Consistently annotating entity mentions in materials science procedural text](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1337–1352, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Alan Ramponi, Barbara Plank, and Rosario Lombardo. 2020. [Cross-domain evaluation of edge detection for biomedical event extraction](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 1982–1989, Marseille, France. European Language Resources Association.
- Mourad Sarrouiti, Carson Tao, and Yoann Mamy Randriamihaja. 2022. [Comparing encoder-only and encoder-decoder transformers for relation extraction from biomedical texts: An empirical study on ten benchmark datasets](#). In *Proceedings of the 21st Workshop on Biomedical Language Processing*, pages 376–382, Dublin, Ireland. Association for Computational Linguistics.
- Daivik Swarup, Ahsaas Bajaj, Sheshera Mysore, Tim O’Gorman, Rajarshi Das, and Andrew McCallum. 2020. [An Instance Level Approach for Shallow Semantic Parsing in Scientific Procedural Text](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3010–3017, Online. Association for Computational Linguistics.
- Shubham Toshniwal, Haoyue Shi, Bowen Shi, Lingyu Gao, Karen Livescu, and Kevin Gimpel. 2020. [A cross-task analysis of text span representations](#). In *Proceedings of the 5th Workshop on Representation Learning for NLP*, pages 166–176, Online. Association for Computational Linguistics.
- Yoshimasa Tsuruoka and Jun’ichi Tsujii. 2005. [Bidirectional inference with the easiest-first strategy for tagging sequence data](#). In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*, pages 467–474, Vancouver, British Columbia, Canada. Association for Computational Linguistics.
- Rob van der Goot. 2021. [We need to talk about train-dev-test splits](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4485–4494, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#).
- Kyosuke Yamaguchi, Ryoji Asahi, and Yutaka Sasaki. 2020. [SC-CoMics: A superconductivity corpus for materials informatics](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 6753–6760, Marseille, France. European Language Resources Association.
- Juntao Yu, Bernd Bohnet, and Massimo Poesio. 2020. [Named entity recognition as dependency parsing](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6470–6476, Online. Association for Computational Linguistics.

Appendix

A Hyperparameters

We use AdamW (Loshchilov and Hutter, 2019) as optimizer for all our models. We use an inverse square-root learning rate scheduler similar to the one used by Vaswani et al. (2017) where ws refers to the number of warmup steps:

$$lr = \sqrt{ws} \cdot \min\left(\frac{1}{\sqrt{step_num}}, step_num \cdot ws^{-1.5}\right)$$

For our measurement identification experiments, we downsample the amount of non-measurement sentences since they represent the majority in the training data. We tune this downsampling rate per model since each BERT variant has been shown to prefer a slightly different one. We apply early stopping after 3 epochs without improvement in terms of F1.

Our NER sequence tagging models are trained with two separate learning rates; one for BERT + linear output layer and another one for the CRF output layer. Both learning rates are reported in the respective column of Table 13. We train for 60–100 epochs, depending on the size of the combined dataset, and take the model with the best evaluation score during this period.

For relation extraction, we use a base learning rate of 4e-5 for all experiments, which we found

Model	LM	LR
CRF	BERT	1e-4/7e-3
	SciBERT	5e-5/7e-3
	MatSciBERT	1e-4/7e-3
Dep. Pars.	BERT	2e-4
	SciBERT	9e-5
	MatSciBERT	3e-4

Table 13: **Hyperparameters:** Learning rates for NER models.

to perform best in preliminary experiments. We employ early stopping with a patience of 15 epochs for all experiments.

Our models are trained with Nvidia A100 and V100 GPUs using the PyTorch framework.

B Details on Biaffine Parser Architecture

We here describe the biaffine parser architecture used to predict relations between named entities. Taking as input the NE embeddings described in Sec. 4.4, head and dependent representations for the i 'th NE are computed via two single-layer feed-forward neural networks:

$$\begin{aligned} \mathbf{h}_i^{head} &= \text{FNN}^{head}(\mathbf{e}_i) \\ \mathbf{h}_i^{dep} &= \text{FNN}^{dep}(\mathbf{e}_i) \end{aligned}$$

These representations are then fed to a biaffine classifier that maps head–dependent pairs onto logit vectors $s_{i,j}$ whose dimensionality corresponds to the inventory of relation labels. Using the softmax operation, these scores are transformed into probability distributions $P(y_{i,j})$ over relation labels:

$$\begin{aligned} \text{Biaff}(\mathbf{x}_1, \mathbf{x}_2) &= \mathbf{x}_1^\top \mathbf{U} \mathbf{x}_2 + W(\mathbf{x}_1 \oplus \mathbf{x}_2) + \mathbf{b} \\ s_{i,j} &= \text{Biaff}(\mathbf{h}_i^{head}, \mathbf{h}_j^{dep}) \\ P(y_{i,j}) &= \text{softmax}(s_{i,j}) \end{aligned}$$

The predicted relation for a pair of named entities is the one receiving the highest probability (which may be \emptyset , i.e., no relation).

Token embeddings. The token embeddings $\mathbf{e}_{i,\text{START}}$ and $\mathbf{e}_{i,\text{END}}$, which form part of the NE embeddings \mathbf{e}_i , are computed as a learned scalar mixture of BERT layers as described by [Kondratyuk and Straka \(2019\)](#).

C Detailed Corpus Statistics

Table 14 shows the NE counts in MuLMS by dataset.

Table 15 and Table 16 show the detailed counts for our inter-annotator agreement study.

Choice of agreement metrics for evaluating agreement on named entity annotations. The task of identifying and labeling NE mentions is a sequence labeling task, hence, κ is not applicable. [Brandesen et al. \(2020\)](#) provide a good explanation of why this is the case in their section 5.1. Using unitizing α_U is an option, but there is no standard implementation or interpretation for NE annotations in the NLP community, and it does not work for overlapping annotations (which we have in our dataset). We opted for using precision and recall, which are intuitively interpretable (*How many of the instances of one type marked by one annotator have also been marked by the respective other annotator?*). [Hripcsak and Rothschild \(2005\)](#) convincingly argue (with a very simple proof) that for sequence labeling tasks such as NEs, F1 actually approaches κ .

Label	total	train	dev	test
MAT	15596	10875	2318	2403
NUM	6081	4142	1077	862
VALUE	4852	3266	895	691
UNIT	4330	2880	789	661
PROPERTY	3925	2867	598	460
FORM	3568	2716	345	507
MEASUREMENT	2171	1531	345	295
CITE	1709	1280	274	155
SAMPLE	1461	1031	249	181
TECHNIQUE	1036	755	146	135
DEV	808	459	235	114
RANGE	736	546	105	85
INSTRUMENT	378	278	50	50

Table 14: Label counts for **named entities** in MuLMS.

D SOFC-Exp and MSPT Corpora

In Sec. 5.4, we perform several multi-task learning (MTL) experiments with MuLMS and two additional NLP datasets in the materials science domain, SOFC-Exp ([Friedrich et al., 2020](#)) and MSTP ([Mysore et al., 2019](#)). We here describe them briefly.

There are 4 named entities in the SOFC-Exp corpus: MATERIAL, which refers to mentions of materials or chemical formulas, VALUE, which denotes numerical values and their corresponding physical unit, DEVICE, which marks device types used in an experiment, and EXPERIMENT, which indicates frame evoking words. Furthermore, there are 16 distinct slots that are modeled as relations between experiment frame evoking word and cor-

Label	P	R	matches		# A1	# A2
			exact	relaxed		
MAT	96.7	91.2	144	145	150	159
NUM	98.9	100.0	86	86	87	86
VALUE	100.0	100.0	53	54	54	54
UNIT	97.9	100.0	47	47	48	47
PROPERTY	42.6	37.7	21	29	68	77
FORM	95.7	86.3	44	44	46	51
MEAS.	44.6	51.0	17	25	56	49
CITE	97.4	97.4	38	38	39	39
SAMPLE	3.1	12.5	1	1	32	8
TECHNIQUE	77.5	59.6	25	31	40	52
DEV	96.0	82.8	18	24	25	29
RANGE	100.0	100.0	24	25	25	25
INSTRUMENT	80.0	76.9	18	20	25	26

Table 15: Inter-annotator agreement: **named entities**. Precision and recall computed from relaxed matches.

Label	P	R	matches	# A1	# A2
propertyValue	81.2	81.2	26	32	32
usedIn	40.0	52.2	12	30	23
hasForm	54.7	71.4	35	64	49
measuresProperty	80.5	72.9	70	87	96
conditionProperty	27.7	59.1	13	47	22
conditionEnvironment	0.0	0.0	0	19	4
usedTogether	13.8	28.6	4	29	14
conditionSampleFeatures	43.1	40.0	22	51	55
usesTechnique	72.9	67.3	35	48	52
takenFrom	33.3	63.6	7	21	11
dopedBy	35.3	50.0	6	17	12
conditionInstrument	24.0	60.0	6	25	10
usedAs	23.5	85.7	12	51	14

Table 16: Inter-annotator agreement for **relations**.

responding entity. Table 17 shows the counts for these relations. These counts are not equal to the ones reported by Friedrich et al. (2020) since we have to remove relations that span across multiple sentences as this case cannot be handled by our relation extraction pipeline.

The MSPT corpus introduces additional the 21 entities:

- PROPERTY-MISC
- PROPERTY-UNIT
- NUMBER
- CHARACTERIZATION-APPARATUS
- APPARATUS-UNIT
- CONDITION-MISC
- META
- SYNTHESIS-APPARATUS
- OPERATION
- AMOUNT-MISC

Label	train	dev	test
<i>AnodeMaterial</i>	220	32	26
<i>CathodeMaterial</i>	173	71	37
<i>Conductivity</i>	36	19	23
<i>CurrentDensity</i>	59	6	17
<i>DegradationRate</i>	15	4	1
<i>Device</i>	311	59	109
<i>ElectrolyteMaterial</i>	187	22	120
<i>FuelUsed</i>	124	28	40
<i>InterlayerMaterial</i>	34	17	6
<i>OpenCircuitVoltage</i>	41	3	25
<i>PowerDensity</i>	138	24	70
<i>Resistance</i>	118	15	57
<i>SupportMaterial</i>	88	13	2
<i>TimeOfOperation</i>	42	3	12
<i>Voltage</i>	30	3	14
<i>WorkingTemperature</i>	330	63	138

Table 17: SOFC-Exp relation counts in our setup.

- AMOUNT-UNIT
- REFERENCE
- PROPERTY-TYPE
- MATERIAL
- MATERIAL-DESCRIPTOR
- APPARATUS-DESCRIPTOR
- APPARATUS-PROPERTY-TYPE
- CONDITION-UNIT
- NONRECIPE-MATERIAL
- CONDITION-TYPE
- BRAND

Table 18 lists the counts of the 14 relations of the MSPT dataset that we use in our MTL experiments.

Label	train	dev	test
<i>Recipe-target</i>	270	53	92
<i>Solvent-material</i>	352	61	107
<i>Atmospheric-material</i>	144	25	35
<i>Recipe-precursor</i>	654	152	199
<i>Participant-material</i>	1315	236	400
<i>Apparatus-of</i>	358	56	93
<i>Condition-of</i>	1378	232	415
<i>Descriptor-of</i>	1157	193	333
<i>Number-of</i>	2114	422	663
<i>Amount-of</i>	1099	244	376
<i>Apparatus-attr-of</i>	66	56	24
<i>Brand-of</i>	326	83	91
<i>Core-of</i>	177	23	89
<i>Next-operation</i>	1311	233	391

Table 18: MSPT relation counts in our setup.

E Detailed Experimental Results

This appendix provides further details on our experimental results. Table 26 depicts the results for identifying sentences containing MEASUREMENT or QUAL_MEAS annotations.

NER. Table 19 and Table 20 report F1 for NER on MuLMS per label. Table 21 gives per-label scores for NER in our MTL experiments. Table 22 and Table 23 provide per-label scores for the SOFC-Exp corpus and MSPT corpus in a single-task setting as well as in a multi-task setting with MuLMS added to the training.

Relation extraction. Table 27 lists per-relation scores when using gold NEs or when using predicted NEs for relation extraction, as well as per-relation scores for the majority baseline. Table 25 shows relation extraction scores per label for both dev and test. Table 24 shows overall results for predicted entities on dev and test.

Label	P	R	F1
MAT	83.6 ±1.3	82.2 ±1.2	82.8 ±0.8
NUM	94.9 ±0.6	94.8 ±1.0	94.9 ±0.7
VALUE	89.4 ±0.6	87.0 ±1.2	88.2 ±0.9
UNIT	94.2 ±0.4	90.4 ±1.1	92.3 ±0.6
PROPERTY	49.8 ±3.0	53.0 ±4.5	51.1 ±1.4
CITE	88.6 ±0.8	87.7 ±2.0	88.2 ±1.3
TECHNIQUE	49.6 ±3.4	51.1 ±5.6	50.1 ±2.9
RANGE	70.3 ±5.8	74.8 ±3.5	72.3 ±3.0
INSTRUMENT	46.7 ±2.7	44.8 ±3.5	45.6 ±2.0
SAMPLE	72.5 ±10.2	36.7 ±6.2	47.9 ±4.1
FORM	66.5 ±3.0	71.4 ±1.9	68.9 ±2.5
DEVICE	82.6 ±2.0	74.9 ±2.5	78.6 ±1.9
MEASUREMENT	61.6 ±2.1	55.3 ±3.0	58.2 ±0.8

Table 19: Per label scores for NER using BILOU tagging and MatSciBERT.

Label	CRF-Tagger	Dep. Parser
MAT	82.8 ±0.8	80.0±0.4
NUM	94.9 ±0.7	94.2±0.4
VALUE	88.2 ±0.9	82.7±1.1
UNIT	92.3 ±0.6	90.8±0.7
PROPERTY	51.1±1.4	51.7 ±2.0
CITE	88.2 ±1.3	85.7±1.5
TECHNIQUE	50.1±2.9	51.4 ±2.4
RANGE	72.3 ±3.0	66.4±4.0
INSTRUMENT	45.6 ±2.0	44.1±3.3
SAMPLE	47.9 ±4.1	29.4±18.1
FORM	68.9 ±2.5	67.6±1.1
DEVICE	78.6 ±1.9	76.4±2.2
MEASUREMENT	58.2±0.8	58.7 ±0.7

Table 20: **Per-Label NER** results on test in terms of **F1** (using MatSciBERT).

Label	Single-Task	+ SOFC	+ MSPT
MAT	82.8	82.3	82.3
NUM	94.9	95.6	95.6
VALUE	88.2	88.1	88.2
UNIT	92.3	92.1	92.6
PROPERTY	51.1	49.4	50.8
CITE	88.2	88.1	87.8
TECHNIQUE	50.1	52.9	50.0
RANGE	72.3	76.2	75.2
INSTRUMENT	45.6	45.9	42.8
SAMPLE	47.9	38.3	45.8
FORM	68.9	69.9	64.9
DEVICE	78.6	77.9	76.3
MEASUREMENT	58.2	57.8	56.9

Table 21: **Per-Label NER** results for MuLMS on test in terms of **F1** for single-task and multi-task MatSciBERT taggers.

Label	Single-Task	+ MuLMS
MATERIAL	75.8	73.2
EXPERIMENT	81.7	81.2
VALUE	93.9	92.0
DEVICE	72.6	72.0

Table 22: **Per-Label Named Entity Recognition** results for SOFC-Exp on test in terms of **F1** using single-task and multi-task MatSciBERT taggers.

Label	ST	+ MuLMS
META	47.5	46.3
PROPERTY-MISC	32.8	34.7
SYNTHESIS-APPARATUS	68.7	66.7
OPERATION	85.0	84.9
PROPERTY-UNIT	42.3	44.5
AMOUNT-MISC	41.4	26.0
NUMBER	94.8	95.5
AMOUNT-UNIT	95.5	95.0
REFERENCE	70.9	67.7
PROPERTY-TYPE	24.6	19.0
MATERIAL	84.1	81.6
MATERIAL-DESCRIPTOR	67.8	63.7
CHARACTERIZATION-APPARATUS	16.2	28.8
APPARATUS-UNIT	57.8	61.5
APPARATUS-DESCRIPTOR	67.0	65.1
APPARATUS-PROPERTY-TYPE	0.0	0.0
CONDITION-MISC	72.3	73.5
CONDITION-UNIT	95.2	94.3
NONRECIPE-MATERIAL	62.3	59.6
CONDITION-TYPE	15.7	12.8
BRAND	71.1	64.0

Table 23: **Per-Label Named Entity Recognition** results for MSPT on test in terms of **F1** using single-task (ST) and multi-task MatSciBERT taggers.

	micro F1	macro F1
dev	42.5 \pm 1.0	37.9 \pm 1.7
test	36.5 \pm 0.9	32.8 \pm 1.2

Table 24: Relation extraction results in terms of **F1**, predicted named entities (including standard deviation over five folds).

Label	dev	test
<i>hasForm</i>	71.3 \pm 0.8	76.1 \pm 0.5
<i>measuresProperty</i>	88.0 \pm 1.0	83.1 \pm 0.8
<i>measuresPropertyValue</i>	84.1 \pm 2.2	73.8 \pm 1.0
<i>usedAs</i>	50.2 \pm 2.3	41.8 \pm 1.9
<i>conditionProperty</i>	83.0 \pm 1.3	72.3 \pm 1.0
<i>conditionPropertyValue</i>	74.7 \pm 2.5	63.2 \pm 2.5
<i>conditionSampleFeatures</i>	67.6 \pm 0.9	66.0 \pm 2.1
<i>usesTechnique</i>	94.6 \pm 0.6	91.0 \pm 0.9
<i>conditionEnvironment</i>	57.1 \pm 6.2	39.0 \pm 3.0
<i>propertyValue</i>	86.7 \pm 1.3	82.5 \pm 2.1
<i>usedIn</i>	49.3 \pm 5.6	37.9 \pm 3.7
<i>conditionInstrument</i>	98.2 \pm 0.8	92.2 \pm 0.9
<i>dopedBy</i>	0.0 \pm 0.0	22.7 \pm 18.7
<i>takenFrom</i>	85.5 \pm 3.7	84.7 \pm 3.6
<i>usedTogether</i>	7.5 \pm 2.0	4.0 \pm 1.7
Macro-avg.	66.5 \pm 1.1	62.0 \pm 1.0
Micro-avg.	73.2 \pm 1.0	67.6 \pm 1.0

Table 25: Per-label F1 scores for relation extraction using MatSciBERT (gold named entities).

LM	Label	dev			test		
		P	R	F1	P	R	F1
BERT	MEASUREMENT	71.5 \pm 1.9	67.1 \pm 3.0	69.2 \pm 1.3	74.1 \pm 3.1	71.4 \pm 5.6	72.5 \pm 2.1
	QUAL_MEAS	45.6 \pm 2.8	61.4 \pm 3.7	52.2 \pm 2.1	49.9 \pm 3.1	51.6 \pm 3.2	50.6 \pm 0.7
	2-Class Macro Avg.	58.5 \pm 2.0	64.3 \pm 2.8	60.7 \pm 1.3	62.0 \pm 3.0	61.5 \pm 3.7	61.5 \pm 1.2
SciBERT	MEASUREMENT	69.1 \pm 1.9	77.6 \pm 1.5	73.1 \pm 0.4	71.1 \pm 2.0	79.5 \pm 1.6	75.0 \pm 0.7
	QUAL_MEAS	54.0 \pm 1.7	66.3 \pm 5.0	59.4 \pm 1.9	52.7 \pm 0.5	52.8 \pm 3.2	52.7 \pm 1.4
	2-Class Macro Avg.	61.5 \pm 1.2	72.0 \pm 2.8	66.2 \pm 1.0	61.9 \pm 0.9	66.2 \pm 1.8	63.9 \pm 0.6
MatSciBERT	MEASUREMENT	69.4 \pm 2.6	77.9 \pm 4.0	73.2 \pm 0.7	70.6 \pm 2.2	80.1 \pm 3.5	74.9 \pm 0.6
	QUAL_MEAS	51.4 \pm 1.4	67.0 \pm 1.6	58.2 \pm 1.1	52.8 \pm 0.8	56.9 \pm 2.6	54.7 \pm 1.0
	2-Class Macro Avg.	60.4 \pm 1.9	72.4 \pm 2.0	65.7 \pm 0.4	61.7 \pm 1.0	68.5 \pm 2.4	64.8 \pm 0.7
<i>human agreement</i>	MEASUREMENT			74.2			74.2
	QUAL_MEAS			61.7			61.7

Table 26: Ternary sentence classification results for identifying sentences containing MEASUREMENT or QUAL_MEAS annotations vs. NONE. Human agreement is only suitable for a rough comparison because it is estimated on a subset of the data.

Label	Gold Entities			Predicted Entities			Majority Baseline		
	P	R	F1	P	R	F1	P	R	F1
<i>hasForm</i>	74.0 \pm 1.8	78.5 \pm 1.1	76.1 \pm 0.5	54.2 \pm 3.6	55.5 \pm 1.6	54.8 \pm 2.3	0.0	0.0	0.0
<i>measuresProperty</i>	80.2 \pm 1.1	86.3 \pm 2.2	83.1 \pm 0.8	39.3 \pm 2.8	36.6 \pm 2.2	37.8 \pm 1.4	50.5	99.6	67.0
<i>measuresPropertyValue</i>	68.6 \pm 1.9	80.0 \pm 1.5	73.8 \pm 1.0	41.5 \pm 4.1	38.3 \pm 4.8	39.7 \pm 3.6	0.0	0.0	0.0
<i>usedAs</i>	49.2 \pm 2.8	36.5 \pm 2.2	41.8 \pm 1.9	30.1 \pm 4.6	19.1 \pm 3.0	23.2 \pm 3.4	0.0	0.0	0.0
<i>conditionProperty</i>	65.3 \pm 2.1	81.2 \pm 2.5	72.3 \pm 1.0	29.0 \pm 3.6	27.4 \pm 2.6	27.9 \pm 1.8	0.0	0.0	0.0
<i>conditionPropertyValue</i>	51.3 \pm 3.0	82.4 \pm 3.3	63.2 \pm 2.5	26.7 \pm 2.8	42.5 \pm 2.7	32.7 \pm 2.4	27.2	100.0	42.7
<i>conditionSampleFeatures</i>	60.7 \pm 2.8	72.3 \pm 1.3	66.0 \pm 2.1	34.4 \pm 4.6	28.4 \pm 3.1	31.0 \pm 3.3	70.4	42.8	53.2
<i>usesTechnique</i>	87.1 \pm 1.1	95.2 \pm 1.3	91.0 \pm 0.9	44.7 \pm 1.6	37.9 \pm 2.7	41.0 \pm 1.8	81.8	100.0	90.0
<i>conditionEnvironment</i>	40.4 \pm 2.6	37.9 \pm 4.1	39.0 \pm 3.0	32.2 \pm 4.1	27.0 \pm 3.8	29.2 \pm 3.1	0.0	0.0	0.0
<i>propertyValue</i>	78.7 \pm 3.1	86.8 \pm 2.7	82.5 \pm 2.1	40.9 \pm 2.6	46.7 \pm 5.1	43.4 \pm 2.0	0.0	0.0	0.0
<i>usedIn</i>	42.5 \pm 4.9	35.6 \pm 7.6	37.9 \pm 3.7	18.6 \pm 7.1	14.2 \pm 3.6	15.9 \pm 4.8	0.0	0.0	0.0
<i>conditionInstrument</i>	93.4 \pm 0.1	91.1 \pm 1.6	92.2 \pm 0.9	38.6 \pm 1.9	36.6 \pm 3.7	37.5 \pm 2.5	90.4	100.0	94.9
<i>dopedBy</i>	26.7 \pm 22.6	20.0 \pm 16.3	22.7 \pm 18.7	20.0 \pm 40.0	6.7 \pm 13.3	10.0 \pm 20.0	0.0	0.0	0.0
<i>takenFrom</i>	75.3 \pm 5.1	96.9 \pm 3.8	84.7 \pm 3.6	67.8 \pm 7.8	61.5 \pm 6.9	64.0 \pm 4.6	46.4	100.0	63.4
<i>usedTogether</i>	9.6 \pm 3.5	2.5 \pm 1.1	4.0 \pm 1.7	9.6 \pm 4.0	2.4 \pm 1.1	3.8 \pm 1.7	0.0	0.0	0.0
Macro-avg.	60.2 \pm 0.8	65.5 \pm 1.6	62.0 \pm 1.0	35.2 \pm 2.7	32.0 \pm 1.6	32.8 \pm 1.2	24.4	36.1	27.4
Micro-avg.	66.8 \pm 1.7	68.4 \pm 0.4	67.6 \pm 1.0	38.6 \pm 2.3	34.7 \pm 1.3	36.5 \pm 0.9	50.5	29.5	37.2

Table 27: Per-label scores (MuLMS test set) for **relation extraction** using MatSciBERT. Majority baseline is computed on gold entities.

Label	MuLMS only			MuLMS + SOFC-Exp			MuLMS + MSPT		
	P	R	F1	P	R	F1	P	R	F1
<i>hasForm</i>	74.0 ±1.8	78.5 ±1.1	76.1 ±0.5	78.0 ±2.9	78.8 ±1.8	78.3 ±0.9	74.2 ±0.7	79.6 ±1.3	76.8 ±0.9
<i>measuresProperty</i>	80.2 ±1.1	86.3 ±2.2	83.1 ±0.8	79.0 ±1.1	87.1 ±0.9	82.8 ±0.4	78.8 ±1.2	86.6 ±1.3	82.5 ±0.3
<i>measuresPropertyValue</i>	68.6 ±1.9	80.0 ±1.5	73.8 ±1.0	68.9 ±3.5	82.7 ±2.2	75.1 ±2.0	70.3 ±0.7	82.1 ±3.1	75.7 ±1.5
<i>usedAs</i>	49.2 ±2.8	36.5 ±2.2	41.8 ±1.9	47.9 ±2.6	35.7 ±2.0	40.8 ±1.3	51.9 ±1.3	37.5 ±1.8	43.5 ±1.4
<i>conditionProperty</i>	65.3 ±2.1	81.2 ±2.5	72.3 ±1.0	67.4 ±2.5	82.1 ±2.5	73.9 ±0.8	66.1 ±1.9	80.3 ±2.2	72.5 ±1.6
<i>conditionPropertyValue</i>	51.3 ±3.0	82.4 ±3.3	63.2 ±2.5	53.7 ±4.2	77.8 ±4.3	63.3 ±2.0	51.1 ±2.6	78.5 ±3.1	61.8 ±2.1
<i>conditionSampleFeatures</i>	60.7 ±2.8	72.3 ±1.3	66.0 ±2.1	61.1 ±3.7	71.4 ±3.8	65.7 ±1.4	62.8 ±3.4	72.2 ±2.5	67.1 ±2.6
<i>usesTechnique</i>	87.1 ±1.1	95.2 ±1.3	91.0 ±0.9	85.6 ±0.8	97.1 ±1.0	91.0 ±0.3	87.2 ±1.1	95.7 ±0.8	91.2 ±0.6
<i>conditionEnvironment</i>	40.4 ±2.6	37.9 ±4.1	39.0 ±3.0	46.3 ±3.5	46.2 ±4.6	46.0 ±2.7	47.2 ±6.3	49.8 ±4.7	48.3 ±5.2
<i>propertyValue</i>	78.7 ±3.1	86.8 ±2.7	82.5 ±2.1	76.2 ±2.7	85.5 ±1.7	80.5 ±1.2	81.4 ±1.7	88.7 ±1.7	84.8 ±1.1
<i>usedIn</i>	42.5 ±4.9	35.6 ±7.6	37.9 ±3.7	45.5 ±7.0	47.6 ±5.7	45.9 ±3.7	41.8 ±4.9	40.4 ±5.5	40.5 ±1.4
<i>conditionInstrument</i>	93.4 ±0.1	91.1 ±1.6	92.2 ±0.9	93.6 ±0.1	93.2 ±0.9	93.4 ±0.5	93.5 ±0.1	91.9 ±1.6	92.7 ±0.9
<i>dopedBy</i>	26.7 ±22.6	20.0 ±16.3	22.7 ±18.7	0.0 ±0.0	0.0 ±0.0	0.0 ±0.0	28.3 ±16.3	26.7 ±13.3	27.0 ±14.0
<i>takenFrom</i>	75.3 ±5.1	96.9 ±3.8	84.7 ±3.6	61.7 ±7.5	98.5 ±3.1	75.6 ±5.7	77.0 ±8.6	98.5 ±3.1	86.2 ±6.2
<i>usedTogether</i>	9.6 ±3.5	2.5 ±1.1	4.0 ±1.7	8.8 ±3.3	2.2 ±0.8	3.5 ±1.2	12.0 ±7.2	3.6 ±3.4	5.4 ±4.8
Macro-avg.	60.2 ±0.8	65.5 ±1.6	62.0 ±1.0	58.2 ±1.5	65.7 ±1.1	61.1 ±0.7	61.6 ±0.9	67.5 ±1.0	63.7 ±0.7
Micro-avg.	66.8 ±1.7	68.4 ±0.4	67.6 ±1.0	67.4 ±1.9	68.9 ±1.2	68.1 ±0.7	68.0 ±1.1	69.5 ±0.5	68.7 ±0.7

Table 28: Per-label scores (MuLMS test set, gold entities) for **multi-task relation extraction**.