

WARM: A Weakly (+Semi) Supervised Math Word Problem Solver

Anonymous ACL submission

Abstract

Solving math word problems (MWP) is an important and challenging problem in natural language processing. Existing approaches to solve MWPs require full supervision in the form of intermediate equations. However, labeling every MWP with its corresponding equations is a time-consuming and expensive task. In order to address this challenge of equation annotation, we propose a weakly supervised model for solving MWPs by requiring only the final answer as supervision. We approach this problem by first learning to generate the equation using the problem description and the final answer, which we subsequently use to train a supervised MWP solver. We propose and compare various weakly supervised techniques to learn to generate equations directly from the problem description and answer. Through extensive experiments, we demonstrate that without using equations for supervision, our approach achieves accuracy gains of 4.5% and 32% over the state-of-the-art weakly supervised approach (Hong et al., 2021), on the standard Math23K (Wang et al., 2017) and AllArith (Roy and Roth, 2017) datasets respectively. Additionally, we curate and release new datasets of roughly 10k MWPs each in English and in Hindi (a low resource language). These datasets are suitable for training weakly supervised models. We also present extension of WARM to semi-supervised learning and present further improvements on results, along with insights.

1 Introduction

A Math Word Problem (MWP) is a numerical problem expressed in natural language (problem description), that can be transformed into an equation (solution expression), which can be solved to obtain the final answer. In Table 1, we present an example MWP. Automatically solving MWPs has recently gained lot of research interest in natural language processing (NLP). The task of automatically solving MWPs is challenging owing to two primary

Problem: It costs Rs 5.0 to buy 10.0 peppermint candies. If the candies all have the same price, how much does it cost to buy 1.0 candy ?
Equation: $X = (5.0 / 10.0) \times 1.0$ (Under full supervision)
Answer: 0.5 (Under weak supervision)

Table 1: Example of a Math Word Problem

reasons: i) *The unavailability of large training datasets with problem descriptions, equations as well as corresponding answers* – as depicted in Table 1, equations can provide **full supervision**, since equations can be solved to obtain the answer, and the answer itself amounts to **weak supervision** only; ii) *Challenges in parsing the problem description and representing it suitably for effective decoding of the equations*. Paucity of completely supervised training data can pose a severe challenge in training MWP solvers. Most existing approaches assume the availability of full supervision in the form of both intermediate equations and answers for training. However, annotating MWPs with equations is an expensive and time consuming task. There exists only two sufficiently large datasets (Wang et al., 2017) in Chinese and (Amini et al., 2019) in English consisting of MWPs with annotated intermediate equations for supervised training.

We propose a novel two-step weakly supervised technique to solve MWPs by making use only of the weak supervision, in the form of answers. In the first step, using only the answer as supervision, we learn to generate equations for questions in the training set. In the second step, we use the generated equations along with answers to train any state-of-the-art supervised model. We illustrate the effectiveness of our weakly supervised approach on our newly curated reasonably large dataset in English and a similarly curated dataset in Hindi - a low resource language. We also perform experiments with semi-supervision and demonstrate how our model can benefit from a small amount of com-

077	pletely labelled data. Our main contributions are	127
078	as follows:	
079	1) An approach, WARM, (<i>c.f.</i> , Section 4) for gener-	128
080	ating equations from MWPs, given (weak) supervi-	129
081	sion only in the form of the final answer.	130
082	2) An extended semi-supervised training method	131
083	to leverage a small amount of annotated equations	132
084	as strong/complete supervision.	133
085	3) A new and relatively large dataset, EW10K, in	134
086	English (with more than 10k instances), for training	135
087	weakly supervised models for solving MWPs (<i>c.f.</i> ,	136
088	Section 3). Given that weak supervision makes it	137
089	possible train MWP solvers even in the absence of	138
090	extensive equation labels, we also present results on	139
091	a similarly crawled dataset, HW10K (with around	140
092	10k instances), in a low resource language, <i>viz.</i>	141
093	Hindi, where we can avoid the additional effort	142
094	required to generate equation annotations.	143
095	4) We empirically show that WARM outperforms	144
096	state-of-the-art models on most of the datasets. Fur-	145
097	ther, we empirically demonstrate the benefits of the	146
098	semi-supervised extension to WARM.	147
099	2 Related Work	148
100	Automatic math word problem solving has recently	149
101	drawn significant interests in the natural language	150
102	processing (NLP) community. Existing MWP solv-	151
103	ing methods can be broadly classified into four	152
104	categories: (a) rule-based methods, (b) statistics-	153
105	based methods, (c) tree-based methods, and (d)	154
106	neural-network-based methods.	155
107	Rule-based systems (Fletcher, 1985; Bakman,	156
108	2007; Yuhui et al., 2010) were amongst the earli-	157
109	est approaches to solve MWPs. They rely heavily	158
110	on hand-engineered rules that might cover a lim-	159
111	ited domain of problems. Statistics-based meth-	160
112	ods (Hosseini et al., 2014; Kushman et al., 2014;	161
113	Sundaram and Khemani, 2015; Mitra and Baral,	162
114	2016; Liang et al., 2016a,b) use predefined logic	163
115	templates and employ traditional machine learning	164
116	models to identify entities, quantities, and oper-	165
117	ators from the problem text and subsequently em-	166
118	ploy simple logical inference to yield the numeric	167
119	answer. (Upadhyay et al., 2016) employ a semi-	168
120	supervised approach by learning to predict tem-	169
121	plates and corresponding alignments using both	170
122	explicit and implicit supervision. Tree-based meth-	171
123	ods (Roy and Roth, 2015; Koncel-Kedziorski et al.,	172
124	2015; Roy et al., 2016; Roy and Roth, 2017, 2018)	173
125	replaced the process of deriving an equation by	174
126	constructing an equivalent tree structure, step by	175
	step, in a bottom-up manner.	176
	More recently, neural network-based MWP solv-	177
	ing methods have been proposed (Wang et al.,	
	2017, 2018a,b; Huang et al., 2018; Chiang and	
	Chen, 2019; Wang et al., 2019; Liu et al., 2019;	
	Xie and Sun, 2019; Wu et al., 2021; Shen et al.,	
	2021). These employ an encoder-decoder archi-	
	tecture and train in an end-to-end manner with-	
	out the need for hand-crafted rules or templates.	
	(Wang et al., 2017) were the first to propose a	
	sequence-to-sequence (Seq2Seq) model, <i>viz.</i> , Deep	
	Neural Solver, for solving MWPs. They employ	
	an RNN-based encoder-decoder architecture to di-	
	rectly translate the problem text into equation tem-	
	plates and also release a high-quality large-scale	
	dataset, Math23K, consisting of 23,161 MWPs in	
	Chinese.	
	(Liu et al., 2019) and (Xie and Sun, 2019) pro-	
	pose tree-structured decoding that generates the	
	syntax tree of the equation in a top-down man-	
	ner. In addition to applying tree-structured decod-	
	ing, (Zhang et al., 2020) propose a graph-based	
	encoder to capture relationships and order informa-	
	tion among the quantities. For a more comprehen-	
	sive review on automatic MWP solvers, readers can	
	refer to a recent survey paper (Zhang et al., 2018).	
	Unlike all the previous works that require equa-	
	tions for supervision, (Hong et al., 2021) propose	
	a weakly supervised method for solving MWPs,	
	where the answer alone is required for training.	
	Their approach attempts to generate the equation	
	tree in a rule based manner so that the correct an-	
	swer is reached. They then train their model using	
	the fixed trees. With the same motivation, we also	
	propose a novel weakly supervised model, WARM,	
	(<i>c.f.</i> , Section 4) for solving MWPs using only the	
	final answer for supervision. We show how WARM	
	can be extended to semi-supervised joint learning	
	in the presence of weak answer-level supervision	
	in conjunction with some equation-level supervi-	
	sion. Further, we empirically demonstrate that WARM	
	outperforms (Hong et al., 2021) on all the datasets.	
	This paper is organized as follows. In Section 3,	
	we set the premise for our approach by describing	
	the new datasets (EW10K and HW10K) for weak	
	supervision that we release. In Section 4, we de-	
	scribe our weakly supervised approach WARM and	
	its semi-supervised extension WARM-S. In Sec-	
	tion 5, we present the experimental setup whereas	
	in Section 6 we delve into the results and its analy-	
	sis before concluding in Section 7.	

178

3 Dataset

179

Currently, there does not exist any sufficiently large English dataset for single and simple equation MWPs. While there exists an English dataset (Amini et al., 2019) with sufficiently large MWPs, the questions in the dataset are meant to be evaluated in a multiple choice question (MCQ) manner. Also, the equation associated with each word problem in this dataset is significantly more complex and requires several binary and unary operators. On the other hand, Math23K (Wang et al., 2017) is in Chinese and Dolphin18k (Huang et al., 2016) contains mostly multi-variable word problems. To address these gaps, we curate a new English MWP dataset, viz., EW10K consisting of 10227 word problem instances (each associated with a single equation) that can be used for training MWP solver models in a weakly supervised manner.

180

We crawled IXL¹ to obtain MWPs for grades VI until X. These word problems involve a wide variety of mathematical computations ranging from simple addition-subtraction to much harder mensuration and probability problems. The dataset consists of 10 different types of problems, spanning 3 tiers of difficulty. We also annotate the dataset with the target unit. The exact distributions are presented in Figure 1.

181

We similarly created a MWP dataset in Hindi - a low resource language. It consists of 9,896 question answer pairs. To the best of our knowledge, this is the first MWP dataset of such size in Hindi.

182

We similarly created a MWP dataset in Hindi - a low resource language. It consists of 9,896 question answer pairs. To the best of our knowledge, this is the first MWP dataset of such size in Hindi.

183

We similarly created a MWP dataset in Hindi - a low resource language. It consists of 9,896 question answer pairs. To the best of our knowledge, this is the first MWP dataset of such size in Hindi.

184

We similarly created a MWP dataset in Hindi - a low resource language. It consists of 9,896 question answer pairs. To the best of our knowledge, this is the first MWP dataset of such size in Hindi.

185

We similarly created a MWP dataset in Hindi - a low resource language. It consists of 9,896 question answer pairs. To the best of our knowledge, this is the first MWP dataset of such size in Hindi.

186

We similarly created a MWP dataset in Hindi - a low resource language. It consists of 9,896 question answer pairs. To the best of our knowledge, this is the first MWP dataset of such size in Hindi.

187

We similarly created a MWP dataset in Hindi - a low resource language. It consists of 9,896 question answer pairs. To the best of our knowledge, this is the first MWP dataset of such size in Hindi.

188

We similarly created a MWP dataset in Hindi - a low resource language. It consists of 9,896 question answer pairs. To the best of our knowledge, this is the first MWP dataset of such size in Hindi.

189

We similarly created a MWP dataset in Hindi - a low resource language. It consists of 9,896 question answer pairs. To the best of our knowledge, this is the first MWP dataset of such size in Hindi.

190

We similarly created a MWP dataset in Hindi - a low resource language. It consists of 9,896 question answer pairs. To the best of our knowledge, this is the first MWP dataset of such size in Hindi.

191

We similarly created a MWP dataset in Hindi - a low resource language. It consists of 9,896 question answer pairs. To the best of our knowledge, this is the first MWP dataset of such size in Hindi.

192

We similarly created a MWP dataset in Hindi - a low resource language. It consists of 9,896 question answer pairs. To the best of our knowledge, this is the first MWP dataset of such size in Hindi.

193

We similarly created a MWP dataset in Hindi - a low resource language. It consists of 9,896 question answer pairs. To the best of our knowledge, this is the first MWP dataset of such size in Hindi.




Figure 1: Distribution on different types of questions

210

4 Our Approach: WARM

211

We propose a weakly supervised model, WARM², for solving the MWP using only the answer for supervision. It is a two-step cascaded approach

212

We propose a weakly supervised model, WARM², for solving the MWP using only the answer for supervision. It is a two-step cascaded approach

213

We propose a weakly supervised model, WARM², for solving the MWP using only the answer for supervision. It is a two-step cascaded approach

¹<https://in.ixl.com/>

²WARM stands for **W**e**A**kly supe**R**vised **M**ath solver.

for weakly supervised MWP solving. For the first step, we propose a model that predicts the equation, given a problem text and answer. This model uses reinforcement learning to search the space of possible equations, given the question and the correct answer only. The answer acts as the goal of the agent and the search is terminated either when the answer is reached or when the equation length exceeds a pre-defined length (this is required, else the search space would be infinitely large). The model is designed to be a two layer bidirectional GRU (Cho et al., 2014) encoder and a decoder network with fully connected units (described in Section 4.3). We refer to this model as WARM. Note that this model requires an answer to determine when to stop exploring. Since we ultimately want a model which should only take the problem statement as input and generate the answer (by generating the correct equation), this model alone is insufficient for evaluation. Using this model, we create a noisy equation-annotated dataset from the weakly annotated training dataset (the training dataset has answers since it is weakly supervised). We use only those instances to create the dataset for which the equation generated by the model yields the correct answer. Note that the equations are noisy, since there is no guarantee that the generated equation will be the shortest or even correct. In the second step, we use this noisy data for supervised training of a state-of-the-art model. The trained supervised model is finally used for evaluation. For simplicity, we provide a summary of notations in Section 1 in supplementary.

214

215

216

217

218

219

220

221

222

223

224

225

226

227

228

229

230

231

232

233

234

235

236

237

238

239

240

241

242

243

244

245

246

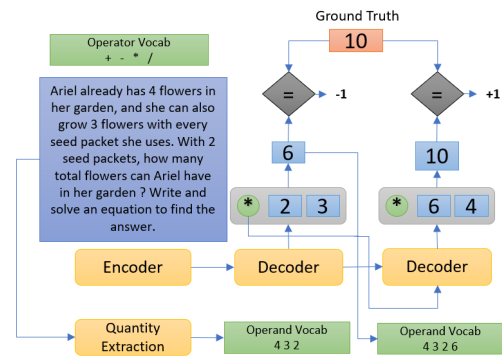


Figure 2: Inference Illustration

4.1 Equation Generation

The first step of our approach is to generate equation given a problem text P and answer A . This is done by using our WARM model. The problem

247

248

249

250

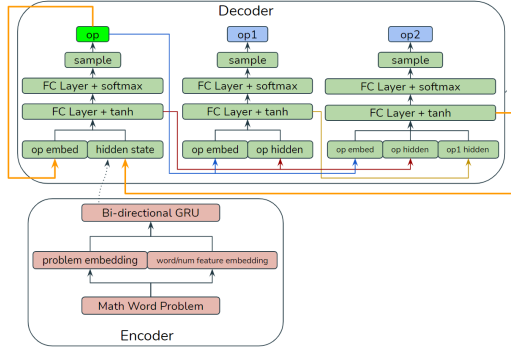


Figure 3: Architecture for generating equation tree in WARM.

text is passed through the encoder of the WARM model to get its encoded representation which is then fed to the decoder. At each time step, the decoder generates an operator and its two operands from the operator and operand vocabulary list. The operation is then executed to obtain a new quantity. This quantity is checked against the ground truth and if it matches the ground truth, the decoding is terminated and a reward of +1 is assigned. Else we assign a reward of -1 and the generated quantity is added to the operand vocabulary list and the decoding continues. The working of the WARM model and architecture are illustrated in Figure 2 and Figure 3 respectively. In the following few subsections, we describe the architecture as well as the training in details.

4.2 Encoder

The encoder takes as input, the MWP represented as a sequence of tokens $P = x_1x_2x_3\dots x_n$. We replace each number in the question with a special token $\langle num_j \rangle$ to obtain this sequence where j denotes the index of number in the operand vocab for that question. Each word token x_i is first transformed into the corresponding word embedding \mathbf{x}_i by looking up an embedding matrix \mathbf{M}_w . Next, a binary feature is appended to the embedding to indicate whether the token is a word or a number. As depicted in the lower half of Figure 3, this appended embedding vector is then passed through a 2 layer bidirectional GRU (Cho et al., 2014) and the outputs from both directions of the final layer are summed to get the encoded representation of the text. This representation is then passed on to the decoder.

4.3 Decoder

The decoder consists of 3 fully connected networks for generating operator, left operand and the right operand. As illustrated in the upper half of Figure 3, the decoder takes as input the previous decoded operand and the last decoder hidden state and outputs the operator, left operand, right operand and hidden state at the current time step. We initialize the decoder hidden state with the last state of the encoder:

$$o_t^p, o_t^l, o_t^r, h_t^d = DecoderFCN(o_{t-1}^p, h_{t-1}^d)$$

Here, h_t^d is the decoder hidden state at the t^{th} time step. o_t^p, o_t^l and o_t^r are probability distributions over operators, left and right operands respectively.

4.3.1 Operator generation

Inside our decoder, we learn an operator embedding matrix $Em_{op}(op_{t-1})$, where op_{t-1} is the operator sampled in the last time step. We generate the operator hidden state h_t^{op} using a gating mechanism.

$$g_t^{op} = \sigma(W_{op}^1[Em_{op}(op_{t-1}); h_{t-1}^d] + b_{op}^1)$$

$$h_t^{op} = g_t^{op} * \tanh(W_{op}^2[Em_{op}(op_{t-1}); h_{t-1}^d] + b_{op}^2)$$

$$o_t^p = \text{softmax}(W_{op}^3 h_t^{op} + b_{op}^3)$$

Here $\sigma()$ denotes the sigmoid function and $*$ denotes elementwise multiplication. We sample operator op_t from the probability distribution o_t^p .

4.3.2 Left Operand Generation

We use the embedding of the current operator $Em(op_t)$ and the operator hidden state h_t^{op} to obtain a probability distribution over the operands. We employ a similar gating mechanism as used for generating operator.

$$g_t^{ol} = \sigma(W_{ol}^1[Em_{op}(op_t); h_t^{op}] + b_{ol}^1)$$

$$h_t^{ol} = g_t^{ol} * \tanh(W_{ol}^2[Em_{op}(op_t); h_t^{op}] + b_{ol}^2)$$

$$o_t^l = \text{softmax}(W_{ol}^3 h_t^{ol} + b_{ol}^3)$$

We sample the left operand ol_t from the probability distribution o_t^l .

4.3.3 Right Operand Generation

For generating the right operand, we use the additional context information that is already available from the generated left operand. Thus, in addition to the operator embedding $Em_{op}(op_t)$ and operator hidden state h_t^{op} we also use the left operand

331 hidden state to get the right operand hidden state
332 h_t^{or} .

$$333 \quad g_t^{or} = \sigma(W_{or}^1[Em_{op}(op_t); h_t^{op}; h_t^{ol}] + b_{or}^1)$$

$$334 \quad h_t^{or} = g_t^{or} * \tanh(W_{or}^2[Em_{op}(op_t); h_t^{op}; h_t^{ol}] + b_{or}^2)$$

$$335 \quad 336 \quad 337 \quad o_t^r = \text{softmax}(W_{or}^3 h_t^{or} + b_{or}^3)$$

338 We sample the right operand or_t from the probabil-
339 ity distribution o_t^r . The hidden state h_t^{or} is returned
340 as the current decoder state h_t^d .

341 4.3.4 Bottom-up Equation Construction

342 For each training instance, we maintain a dictio-
343 nary of possible operands $OpDict$. Initially, this
344 dictionary contains the numeric values from the
345 instance, *i.e.*, the number tokens we have replaced
346 with $\langle num_j \rangle$ during encoding. At the t^{th}
347 decoding step, we sample an operator op_t , left
348 operand ol_t and right operand or_t . We get an inter-
349 mediate result by using the operator corresponding
350 to op_t on the operands ol_t and or_t . This interme-
351 diate result is added to $OpDict$ which enables us to
352 reuse the results of previous computations in future
353 decoding steps. Thus, $OpDict$ acts as a dynamic
354 dictionary of operands and we use it to progress
355 towards the final answer in a bottom-up manner.

356 4.4 Rewards and Loss

357 We use the REINFORCE (Williams, 1992) algo-
358 rithm for training the model using just the final
359 answer as the ground truth. We model the reward
360 as +1 if the predicted answer matches the ground
361 truth and -1 if the predicted answer does not equal
362 the ground truth.

363 Let R_t be defined as the reward obtained after
364 generating $y_t = (op_t, ol_t, or_t)$. The probability P_t
365 of generating the tuple y_t is specified by $p_\theta(y_t) =$
366 $\prod_{i=1}^t o_i^p \times o_i^l \times o_i^r$. The loss is specified as $L =$
367 $-\sum_i \mathbb{E}_{p_\theta(y_i)}[R_i]$ and the corresponding gradient is
368 $\nabla_\theta L = \sum_i \sum_{y_i} p_\theta(y_i) R_i \nabla_\theta \log p_\theta(y_i)$.

369 Since the space of y_i makes it infeasible to com-
370 pute the exact gradient, we use the standardized
371 technique of sampling y_i from $p_\theta(y_i)$ to obtain an
372 estimate of the gradient.

373 4.5 Beam Exploration in Training

374 Since the reward space for our problem is very
375 sparse, we observe that during model training, the
376 gradients go to zero. Our model converges too
377 quickly to some local optima and consequently,

378 the training accuracy saturates to some fixed value
379 despite performing training for a large number of
380 epochs. In order to counter this problem, we em-
381 ploy beam exploration in the training procedure.
382 Instead of sampling operator op_t , left operand ol_t
383 and right operand or_t only once in each decoding
384 step, we sample w triplets (op_t, ol_t, or_t) without re-
385 placement from the joint probability space in each
386 decoding step. Here w is the beam width. This
387 helps in exploring w different paths each epoch,
388 thus increasing the exploration capabilities and re-
389 duce the problem of cold start. In order to select
390 beams from all possible candidates, we have tried
391 multiple heuristics by inspecting the probability
392 and reward values. We have observed empirically
393 that selecting the beam that gives a positive reward
394 at the earliest decoding step yields the best perfor-
395 mance. This enables our model to explore more
396 and mitigates the above problem significantly.

397 4.6 WARM-S: Adding Semi-supervision

398 While it is expensive to completely label large
399 MWP datasets with equations, it is relatively easier
400 to annotate a small percentage of that data. We
401 argue that addition of this small amount of semi-
402 supervision can improve the model training signifi-
403 cantly.

404 We, therefore, consider a model that benefits
405 from a relatively small amount of *strong* supervi-
406 sion in the form of equation annotated data: D_s ,
407 in addition to a potentially larger sized math prob-
408 lem datasets with only *weak* supervision D_w . For
409 a data instance d : $d.p$, $d.e$, and $d.a$ represent
410 its problem statement, equation, and answer respec-
411 tively. D_s consists of instances of the form
412 $(d.p, d.e, d.a)$ while D_w contains instances of the
413 form $(d.p, d.a)$. We extend the WARM model to
414 include a Cross-Entropy loss component for in-
415 stances belonging to D_s . The net loss is the sum of
416 the REINFORCE (RL_{WARM}) and Cross-Entropy
417 losses shown below:-

$$418 \quad \text{Loss 1: } \sum_{d \in D_w} RL_{WARM}(d.p, d.a)$$

$$419 \quad \text{Loss 2: } \sum_{d \in D_s} Cross_Entropy(d.e, WARM(d.p, d.a))$$

420 Thus, we facilitate semi-supervision through
421 **Loss 2**. That is, we jointly use the equations pre-
422 dicted (by WARM) for datapoints belonging to D_w
423 and the ground truth equations for instances be-
424 longing to D_s , for training any state-of-the-art su-
425 pervised MWP solver.

426	5 Experimental Setup	
427	In this section, we report details of the experiments	475
428	on four datasets to examine the performance of the	476
429	proposed weakly supervised model WARM and its	477
430	semi-supervised extension WARM-S. We present	478
431	comparisons with various baselines as well as with	479
432	fully supervised models.	480
433	5.1 Datasets	481
434	We perform all our experiments on the publicly	482
435	available AllArith (Roy and Roth, 2017) and	483
436	Math23K (Wang et al., 2017) datasets and also on	484
437	our EW10K and HW10K datasets. For each dataset,	485
438	we have used a 80 : 20 train-test split.	486
439	AllArith contains 831 MWPs, annotated with equa-	487
440	tions and answers. It is populated by collecting	488
441	problems from smaller datasets, <i>viz.</i> , A12 (Hosseini	490
442	et al., 2014), IL (Roy and Roth, 2015), CC (Roy	491
443	and Roth, 2015) and SingleEQ (Koncel-Kedziorski	492
444	et al., 2015). All mentions of quantities are nor-	493
445	malized to digits. Further, near-duplicate problems	494
446	(with over 80% match of unigrams and bigrams)	495
447	are filtered out.	496
448	Math23K (Wang et al., 2017) contains 23,161	497
449	MWPs in Chinese with 2187 templates, annotated	498
450	with equations and answers, for elementary school	499
451	students and is crawled from multiple online edu-	500
452	cation websites. It is the largest publicly available	501
453	dataset for the task of automatic MWP solving.	502
454	EW10K (<i>c.f.</i> , Section 3) contains 10,227 MWPs	503
455	in English and HW10K contains 9,896 in Hindi	504
456	for classes VI to X. We employ a 80 : 20 train-test	505
457	split in each case.	506
458	5.2 Dataset Preprocessing	507
459	We replace every number token in the problem text	508
460	with a special word token $\langle num_j \rangle$ before pro-	509
461	viding it as input to the encoder. We also define	510
462	a set of numerical constants V_{const} to solve those	511
463	problems which might require special numeric val-	512
464	ues that may not be present in the problem text.	513
465	For example, consider the problem “The radius of	514
466	a circle is 2.5, what is its area?”, the solution is	515
467	“ $\pi \times 2.5 \times 2.5$ ”, but the constant quantity π cannot	516
468	be found in the text. As our model does not use	517
469	equations as supervision, we cannot know precisely	518
470	what extra numeric values might be required for a	519
471	problem, so we fix $V_{const} = \{1, \pi\}$. Finally, the	520
472	operand dictionary for every problem is initialised	521
473	as $OpDict = n_P \cup V_{const}$ where n_P is the set of	522
474	numeric values present in the problem text.	523
	5.3 Implementation Details	
	We implement ³ all our models in PyTorch (Paszke	
	et al., 2019). We set the dimension of the word	
	embedding layer to 128, and the dimension of	
	the hidden states for other layers to 512. We use	
	the REINFORCE (Williams, 1992) algorithm and	
	Adam (Kingma and Ba, 2014) to optimize the pa-	
	rameters. The initial value of the learning rate is	
	set to 0.001, and the learning rate is multiplied by	
	0.7 every 75 epochs. We also set the dropout prob-	
	ability to 0.5 and weight decay to 1e-5 to avoid	
	over-fitting. Finally, we set the beam width to 5	
	in beam exploration. We train our model for 200	
	epochs with the batch size set to 256.	
	5.4 Models	
	We compare the MWP solving accuracy of our	
	weakly supervised models with beam exploration	
	on the following set of baseline and fully super-	
	vised models:	
	WARM is the proposed weakly supervised ap-	
	proach to equation generation (described from Sec-	
	tion 4.1 until 4.4) by employing beam exploration	
	(<i>c.f.</i> , Section 4.5). WARM w/o Beam Exploration	
	is WARM without beam exploration while decod-	
	ing.	
	WARM-S is the semi-supervised extension to	
	WARM (<i>c.f.</i> , Section 4.6) using beam exploration	
	(Section 4.5).	
	WARM-S w/o Beam Exploration is the same as	
	WARM-S but does not use beam exploration while	
	decoding.	
	Random Equation Sampling consists of a ran-	
	dom search over k parallel paths of length d . For	
	each path, an operator and its two operands are	
	uniformly sampled from the given vocabulary and	
	the result is added to the operand vocabulary (sim-	
	ilar to WARM). The equation is terminated once	
	the correct answer is reached. We set $k = 5$ and	
	$d = 40$ for a fair comparison with our model in	
	terms of the number of search operations.	
	Seq2Seq Baseline is a GRU (Cho et al., 2014)	
	based seq2seq encoder-decoder model. REIN-	
	FORCE (Williams, 1992) is used to train the model.	
	Beam exploration is also employed to mitigate is-	
	ssues mentioned in Section 4.5.	
	LBF (Hong et al., 2021) is a weakly supervised	
	model which uses only answer as supervision by	
	fixing incorrect equation parse trees in each iter-	
	ation. It subsequently performs training with the	
	³ Source code is attached as supplementary material	

fixed trees.

Hybrid model w/ SNI (Wang et al., 2017) is a combination of the retrieval and the RNN-based Seq2Seq models with significant number identification (SNI).

Ensemble model w/ EN (Wang et al., 2018a) is an ensemble model that selects the result according to generation probability across Bi-LSTM, ConvS2S and Transformer Seq2Seq models with equation normalization (EN).

Semantically-Aligned (Chiang and Chen, 2019) is a Seq2Seq model with an encoder designed to understand the semantics of the problem text and a decoder equipped with a stack to facilitate tracking the semantic meanings of the operands.

T-RNN + Retrieval (Wang et al., 2019) is a combination of the retrieval model and the T-RNN model comprising a structure prediction module that predicts the template with unknown operators and an answer generation module that predicts the operators.

Seq2Tree (Liu et al., 2019) is a Seq2Tree model with a Bi-LSTM encoder and a top-down hierarchical tree-structured decoder consisting of an LSTM that makes use of the parent and sibling information fed as the input.

GTS (Xie and Sun, 2019) is a tree-structured neural model that generates the expression tree in a goal-driven manner.

Graph2Tree (Zhang et al., 2020) consists of a graph-based encoder which captures the relationships and order information among the quantities. It also employs a tree-based decoder that generates the expression tree in a goal-driven manner.

As described earlier in Section 4, we use our weakly supervised models (WARM and WARM-S) to generate labelled data (*i.e.*, equations) which we then use to train a supervised model. We have performed experiments using GTS (Xie and Sun, 2019) and Graph2Tree (Zhang et al., 2020) as the supervised models since they are the current state-of-the-art.

6 Results and Analysis

6.1 Analyzing WARM

In Table 2, we observe that our model WARM yields far higher accuracy than random baselines with the accuracy values close to 100% on AllArith and EW10K. Thus we are able to more accurately generate equations for a given problem and answer which can then be used to train supervised models.

Weakly Supervised Models	AllArith	Math23K	EW10K	HW10K
WARM w/o Beam Exploration	42.1	14.5	57.5	67.3
WARM	97.4	93.8	99.3	99.5
Baselines	AllArith	Math23K	EW10K	HW10K
Random Equation Sampling	53.4	17.6	46.3	66.6
Seq2Seq Baseline	67.0	7.1	77.6	75.8

Table 2: Equation generation accuracies of WARM based models compared to baselines. All models are trained using ground truth answers on the training set. WARM outperforms all the remaining models by a significant margin on all the datasets. Evidently, beam exploration significantly improves performance.

Weakly Supervised Models	AllArith	Math23K	EW10K	HW10K
WARM w/o Beam Exploration(GTS)	36.1	12.8	52.6	54.1
WARM (GTS)	66.9	55.3	86.9	81.5
WARM w/o Beam Exploration(Graph2Tree)	48.2	13.5	49.8	58.3
WARM (Graph2Tree)	68.7	56.0	87.2	82.9
LBF ‡	51.8	53.6	81.3	75.8
Fully Supervised Models	AllArith	Math23K	EW10K	HW10K
Graph2Tree†	71.9	75.5	NA	NA
GTS†	70.5	73.6	NA	NA
Seq2Tree	–	69.0	NA	NA
T-RNN + Retrieval	–	68.7	NA	NA
Semantically-Aligned†	–	65.8	NA	NA
Ensemble model w/ EN	–	68.4	NA	NA
Hybrid model w/ SNI†	–	64.7	NA	NA

Table 3: MWP solving accuracy of WARM-based models compared to various supervised models on AllArith and Math23K datasets. † denotes that result was reported on 5-fold cross validation. All other models are tested on the test set. ‡ denotes that the result is on the same train-test split as ours. “–” denotes code unavailability/reproducibility issues. NA is not applicable.

Please note that, in Table 2, we report equation generation accuracies on the training set by training the weakly supervised and baseline models using ground truth answers on the training set.

As has been discussed earlier in Section 4.5, our model WARM w/o Beam Exploration suffers from the problem of converging to local optima because of the sparsity of the reward signal. Training our weakly supervised models with beam exploration alleviates the issue to a large extent as we explore the solution space much more extensively and thus

Problem: Ariel already has 4.0 flowers in her garden, and she can also grow 3.0 flowers with every seed packet she uses. With 2.0 seed packets, how many total flowers can Ariel have in her garden ?
Answer: 10.0
Equation Generated: $X=(4.0+(2.0*3.0))$ (Correct)
Problem: Celine took a total of 6.0 quizzes over the course of 3.0 weeks. After attending 7.0 weeks of school this quarter, how many quizzes will Celine have taken in total ? Assume the relationship is directly proportional.
Answer: 14.0
Equation Generated: $X=(7.0+7.0)$ (Incorrect)

Table 4: Equation Generated by WARM model

Problem: Latrell ordered a set of yellow and purple pins. He received 72.0 yellow pins and 8.0 purple pins. What percentage of the pins were yellow?
Equation Generated by WARM (G2T): $X=(72.0*(100.0/(72.0+8.0)))(\text{Correct})$
Equation Generated by LBF: $X=(1.0+(1.0+72.0))$ (Incorrect)
Problem: A square barn has a perimeter of 28.0 metres. How long is each side of the barn ?
Equation Generated by WARM(G2T): $X=((28.0/2.0)/2.0)$ (Correct)
Equation Generated by LBF: $X=((28.0+28.0)/28.0)$ (Incorrect)

Table 5: Comparing WARM and LBF model predicted equations

partly circumventing the sparsity issue. We observe vast improvement in the training accuracy by introduction of beam exploration. The model WARM yields training accuracy significantly higher than its non-beam-explore counterpart. WARM yields the best training accuracy overall. Since the equation generation accuracies of the baselines reported in Table 2 are far worse, the MWP solving accuracies turn out to be significantly worse - around 8-10%, and hence we do not report them.

We also observe that WARM yields results comparable to the various supervised models without requiring any supervision from gold equations. On AllArith, WARM achieves an accuracy of 66.9% and 68.7% using GTS and Graph2Tree as the supervised models respectively. The state-of-the-art supervised model Graph2Tree yields 71.9%. On Math23k, the difference between WARM and the supervised models is more pronounced. WARM’s performance is comparable to that of LBF on Math23k but significantly better on AllArith, EW10K and HW10K, as evident in Table 3

In Table 4, we present some predictions. As can be seen, the model is capable of producing long complex equations as well. Sometimes, it may reach the correct answer but through an incorrect equation. *E.g.:* In the last example, the correct equation would have been $X = 7.0 * 6.0 / 3.0$, but the model predicted $X = 7.0 + 7.0$.

6.2 Analysing Semi-supervision through WARM-S

For analyzing semi-supervision, we combined AllArith (831) with EW10K (10227). We randomly sampled 80% of this data (8846) as our train-set. In retrospect, our train-set consists of 560 instances from AllArith that are completely labelled (amount-

ing to 6.3% of the train-set). We compare our semi-supervised approach against the weakly supervised approach, wherein the entire training data is treated as having only answer labels.

In Table 6 we observe that with less than 10% of fully annotated data, our equation exploration accuracy increases from 56.7% to 92.0% without beam exploration and 99.0% to 99.2% with beam exploration. We also observe a similar trend while training the supervised models; our final MWP solving accuracy increases from 51.2% to 87.4% for WARM w/o Beam Exploration and Graph2Tree as the supervised model. We also study the effect of varying amount of complete supervision in Supplementary Section:2.

Weakly Supervised Models	AllArith +EW10K
WARM w/o Beam Exploration	56.7
WARM	99.0
Semi Supervised Models	AllArith+EW10K
WARM-S w/o Beam Exploration	92.0
WARM-S	99.2

Table 6: Equation generation accuracy of WARM-S compared to weakly supervised models and baselines.

Weakly Supervised Models	AllArith +EW10K
WARM w/o Beam Exploration(GTS)	50.2
WARM (GTS)	87.2
WARM w/o Beam Exploration(Graph2Tree)	51.2
WARM (Graph2Tree)	87.8
Semi Supervised Models	AllArith+EW10K
WARM-S w/o Beam Exploration(GTS)	87.2
WARM-S (GTS)	92.1
WARM-S w/o Beam Exploration(Graph2Tree)	87.4
WARM-S (Graph2Tree)	93.6

Table 7: MWP solving accuracy of WARM-S compared to WARM. With semi-supervision, there is a significant increase in accuracy for WARM w/o Beam Exploration, bringing its performance closer to WARM.

7 Conclusion

We have proposed a two step approach to solving math word problems, using only the final answer for supervision. Our weakly supervised approach, WARM, achieves a reasonable accuracy of 56.0 on the standard Math23K dataset even without leveraging equations for supervision. We also curate and release large scale MWP datasets, EW10K, in English and HW10K, in Hindi. We observed that the results are encouraging for simpler MWPs. We also present the benefits of incorporating a semi-supervised extension to WARM.

649	References	
650	Aida Amini, Saadia Gabriel, Peter Lin, Rik Koncel-	
651	Kedzioriski, Yejin Choi, and Hannaneh Hajishirzi.	
652	2019. Mathqa: Towards interpretable math word	
653	problem solving with operation-based formalisms.	
654	arXiv preprint arXiv:1905.13319 .	
655	Yefim Bakman. 2007. Robust understanding of	
656	word problems with extraneous information. arXiv	
657	preprint math/0701393 .	
658	Ting-Rui Chiang and Yun-Nung Chen. 2019.	
659	Semantically-aligned equation generation for	
660	solving and reasoning math word problems.	
661	In Proceedings of the 2019 Conference of the	
662	North American Chapter of the Association for	
663	Computational Linguistics: Human Language	
664	Technologies . ACL.	
665	Kyunghyun Cho, Bart van Merriënboer, Caglar Gul-	
666	cehre, Dzmitry Bahdanau, Fethi Bougares, Holger	
667	Schwenk, and Yoshua Bengio. 2014. Learning	
668	phrase representations using RNN encoder-decoder	
669	for statistical machine translation . In Proceedings	
670	of the 2014 Conference on Empirical Methods	
671	in Natural Language Processing (EMNLP) , pages	
672	1724–1734, Doha, Qatar. Association for Computa-	
673	tional Linguistics.	
674	Charles R. Fletcher. 1985. Understanding and solving	
675	arithmetic word problems: A computer simulation.	
676	Behavior Research Methods , 17:565–571.	
677	Yining Hong, Qing Li, Daniel Cio, Siyuan Haung, and	
678	Song-Chun Zhu. 2021. Learning by fixing: Solving	
679	math word problems with weak supervision .	
680	Mohammad Javad Hosseini, Hannaneh Hajishirzi,	
681	Oren Etzioni, and Nate Kushman. 2014. Learning	
682	to solve arithmetic word problems with verb catego-	
683	rization. In Proceedings of the 2014 Conference on	
684	Empirical Methods in Natural Language Processing	
685	(EMNLP) , pages 523–533. ACL.	
686	Danqing Huang, Jing Liu, Chin-Yew Lin, and Jian	
687	Yin. 2018. Neural math word problem solver	
688	with reinforcement learning. In Proceedings of	
689	the 27th International Conference on Computational	
690	Linguistics , pages 213–223. ACL.	
691	Danqing Huang, Shuming Shi, Chin-Yew Lin, Jian	
692	Yin, and Wei-Ying Ma. 2016. How well do com-	
693	puters solve math word problems? large-scale	
694	dataset construction and evaluation . In Proceedings	
695	of the 54th Annual Meeting of the Association	
696	for Computational Linguistics (Volume 1: Long	
697	Papers) , pages 887–896, Berlin, Germany. Associ-	
698	ation for Computational Linguistics.	
699	Diederik Kingma and Jimmy Ba. 2014. Adam: A	
700	method for stochastic optimization. International	
701	Conference on Learning Representations .	
702	Rik Koncel-Kedzioriski, Hannaneh Hajishirzi, Ashish	
703	Sabharwal, Oren Etzioni, and Siena Dumas Ang.	
	2015. Parsing algebraic word problems into	704
	equations. Transactions of the Association for	705
	Computational Linguistics , 3:585–597.	706
	Nate Kushman, Yoav Artzi, Luke Zettlemoyer, and	707
	Regina Barzilay. 2014. Learning to automatically	708
	solve algebra word problems. In Proceedings	709
	of the 52nd Annual Meeting of the Association	710
	for Computational Linguistics (Volume 1: Long	711
	Papers) , pages 271–281. Association for Computa-	712
	tional Linguistics.	713
	Chao-Chun Liang, Kuang-Yi Hsu, Chien-Tsung	714
	Huang, Chung-Min Li, Shen-Yu Miao, and Keh-Yih	715
	Su. 2016a. A tag-based english math word problem	716
	solver with understanding, reasoning and explana-	717
	tion. In Proceedings of the Demonstrations Session,	718
	NAACL HLT 2016 , pages 67–71. The Association	719
	for Computational Linguistics.	720
	Chao-Chun Liang, Kuang-Yi Hsu, Chien-Tsung	721
	Huang, Chung-Min Li, Shen-Yu Miao, and Keh-	722
	Yih Su. 2016b. A tag-based statistical english	723
	math word problem solver with understanding,	724
	reasoning and explanation. In Proceedings of	725
	the Twenty-Fifth International Joint Conference on	726
	Artificial Intelligence, IJCAI’16 , page 4254–4255.	727
	AAAI Press.	728
	Qianying Liu, Wenyv Guan, Sujian Li, and Daisuke	729
	Kawahara. 2019. Tree-structured decoding for solv-	730
	ing math word problems. In Proceedings of the	731
	2019 Conference on Empirical Methods in Natural	732
	Language Processing and the 9th International	733
	Joint Conference on Natural Language Processing	734
	(EMNLP-IJCNLP) , pages 2370–2379.	735
	Arindam Mitra and Chitta Baral. 2016. Learning to	736
	use formulas to solve simple arithmetic problems.	737
	In Proceedings of the 54th Annual Meeting of the	738
	Association for Computational Linguistics (Volume	739
	1: Long Papers) , pages 2144–2153. Association for	740
	Computational Linguistics.	741
	Adam Paszke, Sam Gross, Francisco Massa, Adam	742
	Lerer, James Bradbury, Gregory Chanan, Trevor	743
	Killeen, Zeming Lin, Natalia Gimelshein, Luca	744
	Antiga, Alban Desmaison, Andreas Köpf, Edward	745
	Yang, Zach DeVito, Martin Raison, Alykhan Tejani,	746
	Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Jun-	747
	jie Bai, and Soumith Chintala. 2019. Pytorch: An	748
	imperative style, high-performance deep learning li-	749
	brary .	750
	Subhro Roy and Dan Roth. 2015. Solving general	751
	arithmetic word problems. In Proceedings of the	752
	2015 Conference on Empirical Methods in Natural	753
	Language Processing , pages 1743–1752. ACL.	754
	Subhro Roy and Dan Roth. 2017. Unit dependency	755
	graph and its application to arithmetic word prob-	756
	lem solving. In Proceedings of the Thirty-First	757
	AAAI Conference on Artificial Intelligence , page	758
	3082–3088. AAAI Press.	759

760	Subhro Roy and Dan Roth. 2018. Mapping to declarative knowledge for word problem solving. <u>Trans. Assoc. Comput. Linguistics</u> , 6:159–172.	<u>Processing (Volume 1: Long Papers)</u> , pages 5859–5869.	816
761			817
762			
763	Subhro Roy, Shyam Upadhyay, and Dan Roth. 2016. Equation parsing : Mapping sentences to grounded equations. In <u>Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016</u> , pages 1088–1097. The Association for Computational Linguistics.	Zhipeng Xie and Shichao Sun. 2019. A goal-driven tree-structured neural model for math word problems. In <u>Proceedings of the 28th International Joint Conference on Artificial Intelligence</u> , pages 5299–5305. AAAI Press.	818
764			819
765			820
766			821
767			822
768			
769			
770	Jianhao Shen, Yichun Yin, Lin Li, Lifeng Shang, Xin Jiang, Ming Zhang, and Qun Liu. 2021. Generate & rank: A multi-task framework for math word problems. <u>arXiv preprint arXiv:2109.03034</u> .	M. Yuhui, Z. Ying, C. Guangzuo, R. Yun, and H. Ronghuai. 2010. Frame-based calculus of solving arithmetic multi-step addition and subtraction word problems. In <u>2010 Second International Workshop on Education Technology and Computer Science</u> , volume 2, pages 476–479.	823
771			824
772			825
773			826
774	Sowmya S Sundaram and Deepak Khemani. 2015. Natural language processing for solving simple word problems. In <u>Proceedings of the 12th International Conference on Natural Language Processing</u> , pages 394–402. NLP Association of India.	Dongxiang Zhang, Lei Wang, Nuo Xu, Bing Tian Dai, and Heng Tao Shen. 2018. The gap of semantic parsing: A survey on automatic math word problem solvers. <u>CoRR</u> , abs/1808.07290.	829
775			830
776			831
777			832
778			
779	Shyam Upadhyay, Ming-Wei Chang, Kai-Wei Chang, and Wen-tau Yih. 2016. Learning from explicit and implicit supervision jointly for algebra word problems. In <u>Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing</u> , pages 297–306.	Jipeng Zhang, Lei Wang, Roy Ka-Wei Lee, Yi Bin, Yan Wang, Jie Shao, and Ee-Peng Lim. 2020. <u>Graph-to-tree learning for solving math word problems</u> . In <u>Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics</u> , pages 3928–3937, Online. Association for Computational Linguistics.	833
780			834
781			835
782			836
783			837
784			838
785			839
786	Lei Wang, Yan Wang, Deng Cai, Dongxiang Zhang, and Xiaojiang Liu. 2018a. Translating math word problem to expression tree. In <u>Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing</u> , pages 1064–1069. ACL.		
787			
788			
789			
790	Lei Wang, Dongxiang Zhang, Lianli Gao, Jingkuan Song, Long Guo, and Heng Tao Shen. 2018b. Mathdqn: Solving arithmetic word problems via deep reinforcement learning. In <u>Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, AAAI</u> , pages 5545–5552. AAAI Press.		
791			
792			
793			
794			
795			
796	Lei Wang, Dongxiang Zhang, Jipeng Zhang, Xing Xu, Lianli Gao, Bing Tian Dai, and Heng Tao Shen. 2019. Template-based math word problem solvers with recursive neural networks. In <u>Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence, AAAI</u> , pages 7144–7151. AAAI Press.		
797			
798			
799			
800			
801			
802	Yan Wang, Xiaojiang Liu, and Shuming Shi. 2017. Deep neural solver for math word problems. In <u>Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing</u> , pages 845–854. ACL.		
803			
804			
805			
806			
807	Ronald J. Williams. 1992. <u>Simple statistical gradient-following algorithms for connectionist reinforcement learning</u> . <u>Mach. Learn.</u> , 8(3–4):229–256.		
808			
809			
810	Qinzhao Wu, Qi Zhang, Zhongyu Wei, and Xuan-Jing Huang. 2021. Math word problem solving with explicit numerical values. In <u>Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language</u>		
811			
812			
813			
814			
815			

Supplementary Material : WARM: A Weakly(+Semi) Supervised Math word Problem Solver

Anonymous ACL submission

001	1 Notations	3 Infrastructre Details	015
002	We summarize the notations used in section 4 of the main paper in table 1.	GPU Model used : 1)Model number: GeForce GTX 1080 Ti 2)Memory : 12GB	016 017 018 019
		Training time : 1) WARM takes 4 hours for training 2) G2T takes 1 hour and 30 minutes to get trained completely	020 021 022 023 024

Notation	Description
o_t^p	Probability distribution of operators at decoding timestep t.
o_t^l	Probability distribution of the left operand at decoding timestep t.
o_t^r	Probability distribution of the right operand at decoding timestep t.
h_t^d	Decoder hidden state at timestep t.
Em_{op}	Operator Embedding Matrix
h_t^{op}	Hidden state for the operator at timestep t
op_t	Operator sampled from o_t^p .
h_t^{ol}	Hidden state for the left operand at timestep t.
ol_t	Left operand sampled from o_t^l .
h_t^{or}	Hidden state for the right operand at timestep t.
or_t	Right operator sampled from o_t^r .
OpDict	Operand dictionary used while decoding
R_t	Rewards obtained at timestep t.
$p_\theta(y_t)$	Probability of generating $y_t = (op_t, ol_t, or_t)$ at timestep t.

Table 1: Summary of notation used.

003

2 Ablation Study: Varying Amount of Semi-supervision

006 We performed an experiment to study the effect of different amounts of supervision by varying the number of instances in training set we treat as fully labelled. The number of fully labelled instances is X-axis*80. We observe that just having 160 equation-labelled instances (out of 8846 ie. 1.8%) improves the equation-exploration accuracy significantly (46.7% to 90.6%) when we don't use beam exploration.




Figure 1: Equation Exploration accuracy with varying supervision