Who Should Do It? Automatic Identification of Responsible Stakeholder in Writings during Training

Martin Ruskov
Department of Languages, Literatures, Cultures and Mediations
Università degli Studi di Milano, Milan, Italy
martin.ruskov@unimi.it

***Abstract*—In online Problem-Based Learning (PBL), being able to provide immediate feedback to learners is invaluable, yet difficult to achieve. We examine how well an off-the-shelf Natural Language Processing (NLP) framework is able to detect the absence of an identified responsible stakeholder in ideas generated during security training. Part-of-Speech Tagging and Dependency Parsing are applied on contextualised written learner contributions, collected from a PBL environment and compare the results to an assessment performed by experts. Using grammatical analysis, we aim to detect the absence of an identified responsible stakeholder in collected contributions ($n = 1174$) from two security domains. Four heuristics are compared, resulting in a precision of ($PPV = 0.929$) on the best of these, sufficient to provide immediate feedback to learners. Our results suggest that for the purposes of scaffolding open-ended PBL exercises, off-the-shelf NLP frameworks can achieve good performance on responsible stakeholder identification.**

***Index Terms*—dependency parser, part-of-speech tagger, formative feedback, problem-based learning, security training**

I. Introduction

In online learning environments, providing immediate feedback is critical to educational efficiency. Yet, it is difficult to achieve in Problem-Based Learning (PBL) where interactions are complex and multiple solutions could be valid. One particular challenge is evaluating whether written proposed interventions by learners during security training are elaborated in enough detail to become actionable.

Consider the issue of repeated security premise intrusions. A natural candidate for a solution would be to aim to tighten access. However, this could be done in different ways as the following real learner contributions illustrate. The abstract idea "coming up with new policies and rules" might take different shapes depending on who implements it. For example "More security to stop hacking. It staff, IT programmers responsible" might lead to employing bespoke solutions, whereas "dedicated data security contractor to manage security" might bring know-how from other contexts. As visible from these examples, the responsible stakeholder has an important role in shaping the solution, and is thus a key part of it. When ideas are put in writing, grammatically the responsible stakeholder is typically – but not always – indicated by the subject.

We study how grammatical analysis performed by an off-the-shelf Natural Language Processing (NLP) pipeline can identify the absence of responsible stakeholder in learner contributions. In particular, we study written security intervention ideas generated with the CCO Toolkit, an online environment for PBL in the domain of security. We experiment with the combination of two NLP pipelines and two simple heuristics as possible mechanisms to provide immediate feedback to learners. The most successful among these is intended to be integrated in the CCO Toolkit. Here we evaluate the approach on real-world contributions generated by learners who previously used the toolkit. The approach does not rely on any contextual training or adaptation of the NLP pipelines. Thus, we expect it to be generalisable to other domains utilising problem-solving via short proposals for interventions.

II. Background

PBL has been widely used as a framework for security training due to its inherent stimulation of active learning [1]. Typically, given specific problems, learners are asked to discuss and exchange experiences and propose ideas. In an online environment, for PBL it is important to allow for unrestrained learner creativity, e.g. via free-form learner input [2, 74-80]. In crime science, Ekblom defines prevention competences to include *know who to involve* [3, 19-24]. Nutley et al have claimed Ekblom's competences to be essential to wider evidence-based practice and have made a case that they contribute to making a difference between knowing and doing [4]. Accordingly, indications of these competences could be sought in learner input in PBL.

However, the openness of learner input raises challenges for automation in assessment and calls for interpretation beyond lexical parsing [5]. In what could be considered to be an attempt to overcome this limitation, Bagaria et al [6] aim to identify and extract a *subject-verb-object* triplet from sentences contributed by learners, but they do not measure the accuracy of the approach. Beyond that, to our knowledge no previous research has worked on using identification of the responsible stakeholder for the purposes of learning assessment.

Broader advancements in NLP offer a grammatical toolset that can address the challenge of providing feedback on unconstrained discussion. In particular dependency parsers and Part-of-Speech (PoS) taggers could allow for partial – minimal, yet sufficient – real-time interpretation of learner contributions. Such approaches were used even before the advent of transformer models that have revolutionised NLP [7]. Papadimitriou et al [8] also demonstrated that transformers capture grammar information in the embedding itself. Recently, grammatical analysis has been applied to information

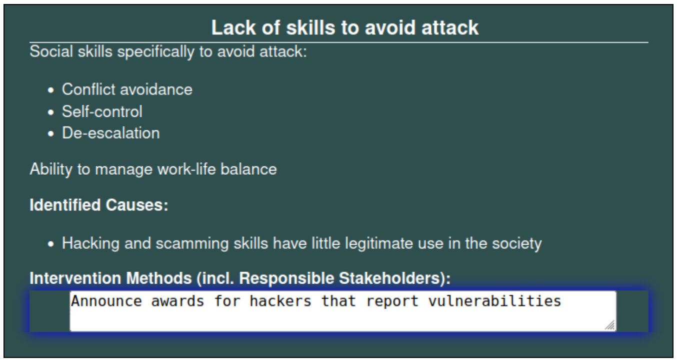

Fig. 1. The interventions prompt in the CCO toolkit with a sample idea. Notice that the prompt explicitly asks for responsible stakeholders.

security reports [9], [10]. However, in short contextualised contributions, like in the case of the CCO Toolkit, there's an additional layer of ambiguity leaving it unclear where context, relevant to the contribution comes from [5].

III. Approach

We analyse a dataset of security intervention ideas provided as short texts by learners using the CCO Toolkit. For these learner contributions, we experiment with heuristics detecting if they do not indicate a responsible stakeholder.

The toolkit is a software platform developed to support the teaching of a related crime prevention theoretical framework [11]. It guides learners through a step-wise brainstorming process with the goal to solve given security problems. The steps of this process are as follows: 1) scenario, 2) identification of causes, 3) identification of potential interventions that could contribute to a solution (see Fig. 1), 4) review of possible interplay between identified interventions, 5) peer-review of intervention ideas by other learners from the perspective of key scenario stakeholders; and 6) feedback and score of learner's own intervention ideas, based on feedback by others and automated analysis. The iterative process described above leads to the inherent complexity of the CCO Toolkit due to combining the underlying theoretical model (represented in the upper part of Fig. 1), the problem scenarios and previous learner contributions (i.e. causes identified in Step 2 and also visible in Fig. 1). This complexity leads to distinct contexts for each contribution, which remain external to its actual content, yet are explicitly or implicitly referenced by it [12].

The analysed dataset[1] includes 1174 intervention contributions by 91 graduate students from two universities. Students worked on three problem scenarios – two from information security ($n = 1041$) and one from community safety ($n = 133$). Working language was English and part of the students were non-native speakers. Intervention contributions lengths range from 4 (see S8 in Table I) to 253 characters (median 58). Two security experts independently annotated interventions indicating the *absence of a responsible stakeholder* (82.7%). This annotation code concerns whether it is clear (e.g. S2, S4) or not (S0, S5) who is enacting the intervention, regardless of grammatical structure. Then the experts were asked to discuss disagreements in their annotations and come up with a compromise. To allow for the interpretation of inaccuracies of the heuristics, the experts were also asked to annotate when contributions were ungrammatical (12.3%) or ambiguous (30.5%).

[1]Full dataset and heuristics code at: https://cco.works/opendata

TABLE I
Samples from the analysed interventions dataset[1].
Key: (S)takeholder absent, (U)ngrammatical and (A)mbiguous.

#	Sample Intervention	Annotation S	U	A
S0	Build trust.			✓
S1	strengthen the security of the information system by hiring professional team to do it. the company management will be responsible for this.			
S2	a dedicated data security contractor to manage security			
S3	More security to stop hacking. It staff, IT programmers responsible		✓	✓
S4	Reduce the accessibility to the back office. (Security Management)		✓	
S5	coming up with new policies and rules	✓		✓
S6	Relook at how resources will be distributed so that more resources will be invested in tackling the problem in the long run - by the management			
S7	Checks to ensure only the most enthusiastic and genuine people are recruited	✓		
S8	CCTV	✓		
S9	administrator/network engineer			

We address the task of identifying the absence of corresponding responsible stakeholder with experiments with two generic NLP pipelines: one based on syntactic word embeddings[2] and one transformer-based[3] We apply PoS tagger and dependency parser on each intervention. Our heuristic algorithms consider candidates for potetial stakeholders from all possible sentences in the contribution. As seen in Table II, candidates are tokens that are 1) `subjects`, but not `pronouns`, or 2) `nouns`, child of an `agent`[1].

To improve accuracy, we further test a simplifying heuristic that performs pruning of the dependency tree. The rationale is that whenever present the responsible stakeholder is unlikely to be in a subordinate clause of a sentence. Our pruning collapses noun branches of the dependency tree that do not contain an agent dependency. For each of the pipelines, this leads us to further experimental setups, one with pruning and one without.

TABLE II
Heuristic criteria for identification of stakeholders according to grammatical voice.

voice	PoS	dependency
active	not `PRON`	`nsubj`
passive	`NOUN`	child of `agent`

IV. Experiments

We compare the results of our heuristics against the baseline of expert annotations. Since we aim to give feedback to learners in case of missing stakeholder, in our classification

[2]https://spacy.io/models/en#en_core_web_lg
[3]https://spacy.io/models#design-trf

TABLE III
COMPARISON OF THE USED HEURISTICS.

Model	PPV	TPR	$F_{0.5}$	bACC
lg	0.924	0.821	0.901	0.748
lg+p	0.923	0.857	0.909	0.758
trf	**0.929**	0.841	0.910	0.768
trf+p	0.927	**0.875**	**0.916**	**0.773**

the positive case denotes this absence. For this reason, precision (PPV) is the count of correctly classified the missing stakeholder over all where the algorithm does not detect it and recall (TPR) is the same value over all interventions where the algorithm classification matches the one by experts. For the same reason, we report a generalised F_β-score emphasising precision ($\beta = 0.5$). Due to small percentage of contributions with a stakeholder present in the dataset, balanced accuracy (bACC) is also reported.

Table III shows the results with respect to the used standard metrics. The transformer-based pipeline (trf) slightly outperform the syntactic one (lg). Overall, we report best achieved results to be precision ($PPV = 0.929$) without pruning (trf) and generalised F-score ($F_{0.5}= 0.916$) with pruning (trf+p). These are excellent results for an untuned off-the-shelf pipeline. Notably, the pruning heuristic brings improvements on all metrics, except precision, but with a very small margin. The performance of our heuristics was consistent across the two application domains (information security and community safety, see Section III).

Whenever contributions contain full phrases with a `subject` the heuristics identifies the responsible stakeholder correctly (e.g. S1 from Table I). When a contribution contains only a `noun` or a noun phrase, it is difficult to interpret if it is a `subject` (S9) or `object` (S8). A related difficulty is present also in more general cases when a `verb` is omitted (S3). There are cases (S4) where punctuation is unambiguous and – despite the missing verb – from the context of the toolkit, a critical interpretation could deduce that the `noun` phrase at the end of the contribution refers to the stakeholder. For the purposes of the toolkit this could be encoded as an explicit rule, but for the sake of generalisability to other domains, we report results without such bespoke logic. In cases when passive voice was used (S6), the `subject` was also correctly identified via the `agent` dependency. Without the pruning heuristic, sometimes the `object` is wrongly identified to be the responsible stakeholder (e.g. “people” in S7). A notable part of the contributions falsely identified as missing information are ungrammatical (50.7%) or ambiguous (43.3%). These are much higher concentrations than the average for the dataset, thus a very probable cause of the errors.

V. FUTURE WORK

These results are encouraging for the CCO Toolkit and for further applications of this approach. An integration of the heuristic in the toolkit is due in order to provide immediate formative feedback to unconstrained input from learners. The presence of classification errors suggests that the automatically-generated feedback should be provided in a non-intrusive manner, making it easy for learners to purposefully ignore. A further usability study is needed to see if this feature actually leads learners to more often identify responsible stakeholders in their contributions and if this benefits the overall PBL experience. Such study would also explore attitudes of learners towards the received feedback in context. Beyond that, the approach could be applied without any adaptations to other problem domains and its transferability should be investigated.

Also, expanding NLP support for stakeholder analysis in the form of entity recognition and matching might help better interpret the practicalities of learner contributions [13]. In particular, this would allow for a better distinction between `subject` and `object` in noun phrases. Also, it is worthwhile to explore the possibility to encourage learners to address other important dimensions of idea quality, such as featured action, responding to another question from Ekblom’s prevention competences – *how to put in practice* [3].

REFERENCES

[1] J. Belur, W. Agnew-Pauley, B. McGinley, and L. Tompson, “A Systematic Review of Police Recruit Training Programmes,” *Policing: A Journal of Policy and Practice*, vol. 14, no. 1, pp. 76–90, 2019.

[2] A. Walker, H. Leary, C. Hmelo-Silver, and P. A. Ertmer, *Essential Readings in Problem-Based Learning: Exploring and Extending the Legacy of Howard S. Barrows*. Purdue University Press, 2015.

[3] P. Ekblom, *Crime Prevention, Security and Community Safety Using the 5Is Framework*, ser. Crime Prevention and Security Management. Palgrave Macmillan UK, 2010.

[4] S. Nutley, I. Walter, and H. T. O. Davies, “From Knowing to Doing: A Framework for Understanding the Evidence-into-Practice Agenda,” *Evaluation*, vol. 9, no. 2, pp. 125–148, 2003.

[5] A. Maya, J. Nazura, and B. L. Muralidhara, “Recent trends in answer script evaluation – a literature survey,” in *Proceedings of the 3rd International Conference on Integrated Intelligent Computing Communication & Security (ICIIC 2021)*. Atlantis Press, 2021, pp. 105–112.

[6] V. Bagaria, M. Badve, M. Beldar, and S. Ghane, “An intelligent system for evaluation of descriptive answers,” in *2020 3rd International Conference on Intelligent Sustainable Systems (ICISS)*, 2020, pp. 19–24.

[7] O. Sychev, A. Anikin, and A. Prokudin, “Automatic grading and hinting in open-ended text questions,” *Cognitive Systems Research*, vol. 59, pp. 264–272, 2020.

[8] I. Papadimitriou, E. A. Chi, R. Futrell, and K. Mahowald, “Deep subjecthood: Higher-order grammatical features in multilingual BERT,” in *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*. Online: Association for Computational Linguistics, Apr. 2021, pp. 2522–2532.

[9] Z. Zhu and T. Dumitras, “Chainsmith: Automatically learning the semantics of malicious campaigns by mining threat intelligence reports,” in *2018 IEEE European Symposium on Security and Privacy (EuroS&P)*, 2018, pp. 458–472.

[10] P. Gao, F. Shao, X. Liu, X. Xiao, Z. Qin, F. Xu, P. Mittal, S. R. Kulkarni, and D. Song, “Enabling efficient cyber threat hunting with cyber threat intelligence,” in *2021 IEEE 37th International Conference on Data Engineering (ICDE)*, 2021, pp. 193–204.

[11] M. Ruskov, J. M. Celdran, P. Ekblom, and M. A. Sasse, “Unlocking the next level of crime prevention: development of a game prototype to teach the conjunction of criminal opportunity,” *Information Technologies and Control*, vol. 10, no. 3, pp. 15–21, 2012.

[12] M. Ruskov, P. Ekblom, and M. A. Sasse, “Getting users smart quick about security: Results from 90 minutes of using a persuasive toolkit for facilitating information security problem solving by non-professionals,” 2022. [Online]. Available: https://doi.org/10.48550/ARXIV.2209.02420

[13] M. S. Reed, A. Graves, N. Dandy, H. Posthumus, K. Hubacek, J. Morris, C. Prell, C. H. Quinn, and L. C. Stringer, “Who’s in and why? A typology of stakeholder analysis methods for natural resource management,” *Journal of Environmental Management*, vol. 90, no. 5, pp. 1933–1949, 2009.