

Entity-Level Sentiment Analysis (ELSA): An exploratory task survey

Egil Rønningstad and Erik Velldal and Lilja Øvrelid

University of Oslo

Department of Informatics

{egilron,erikve,liljao}@ifi.uio.no

Abstract

This paper explores the task of identifying the overall sentiment expressed towards volitional entities (persons and organizations) in a document – what we refer to as Entity-Level Sentiment Analysis (ELSA). While identifying sentiment conveyed towards an entity is well researched for shorter texts like tweets, we find little to no research on this specific task for longer texts with multiple mentions and opinions towards the same entity. This lack of research would be understandable if ELSA can be derived from existing tasks and models. To assess this, we annotate a set of professional reviews for their overall sentiment towards each volitional entity in the text. We sample from data already annotated for document-level, sentence-level, and target-level sentiment in a multi-domain review corpus, and our results indicate that there is no single proxy task that provides this overall sentiment we seek for the entities at a satisfactory level of performance. We present a suite of experiments aiming to assess the contribution towards ELSA provided by document-, sentence-, and target-level sentiment analysis, and provide a discussion of their shortcomings. We show that sentiment in our dataset is expressed not only with an entity mention as target, but also towards targets with a sentiment-relevant relation to a volitional entity. In our data, these relations extend beyond anaphoric coreference resolution, and our findings call for further research of the topic. Finally, we also present a survey of previous relevant work.

1 Introduction

Over the course of the last two decades, the field of NLP has generated a vast body of research on sentiment analysis (SA), i.e. the task of identifying opinions expressed in text. Prior work has focused on a range of different levels of analysis; from document- or sentence-level polarity classification to more fine-grained prediction of various components of opinions, like source/holder expressions,

polar expressions, target expressions and aspect-based sentiment classification.

However, we observe that a more aggregated level of analysis, what we here dub "entity-level sentiment analysis" (or ELSA for short), remains under-explored. For our purposes, we will define ELSA as the task of determining the overall (i.e. document-level) polarity (positive/negative) expressed towards an entity in a text. Moreover, for the current paper we will restrict the discussion to volitional entities, like persons (PER) or organizations (ORG).

A given text might make reference to multiple distinct entities, each of which might be mentioned multiple times, both directly and indirectly, and also have multiple opinions directed towards them. Hence, solving the task of ELSA in its full complexity may potentially involve several different sub-tasks, including (but not necessarily limited to) named entity recognition (NER), resolution of entity mentions, coreference- and anaphora-resolution, identification of sentiment targets and/or aspects and their polarities, and finally what we here refer to as "target–entity resolution", i.e. the task of identifying the particular entity with which a given target expression is associated.

The main goal of this paper is to shed more light on this task of entity-level sentiment analysis. To better understand the complexity of the task, we quantify how far we can potentially get toward the goal of ELSA by simply building on existing tools for SA at lower levels. We do this through an exploratory analysis of an existing SA dataset that comprises annotations for several levels of granularity, for which we experiment with different strategies for aggregating these gold annotations to infer entity-level sentiment. Importantly, we also discuss which pieces appear to currently be missing in order to fully solve ELSA. We start, however, by surveying relevant prior work, and also discuss the often diverging terminology used in the field.

The paper is structured as follows. In Section 2 we first present background literature on the task of analyzing sentiment towards individual entity representations, before we discuss the limited work we found on resolving sentiment-relevant elements to the document-level. We found no previous work describing the ELSA task as we defined it, and very little in general on resolving sentiment from the sub-sentence level to the document-level. Section 3 presents the datasets we sample from; Norwegian data annotated for document-level, sentence-level and target-level sentiment classification. We also present our new, exploratory dataset, annotated directly for document-level sentiment classification for each volitional entity in the text. In Section 4 we present our results from attempting to derive ELSA from document-level and sentence-level sentiment analysis – what we dub target-independent approaches. We find that merely locating an entity mention inside a positive document is a very weak indicator of a positive sentiment towards that entity. Section 5 presents our findings from deriving ELSA classification from Targeted Sentiment Analysis (TSA). Besides discrepancies from annotator disagreement, we find that in order to fully solve the ELSA task, future work should aim to add another level of analysis, corresponding to the relations between sentiment targets and their affiliated entities. In Section 6 we report on a baseline model for ELSA through TSA as a proxy task.

2 Literature review

This section first presents prior work related to ELSA. We then survey the terminology used in the literature, comparing this to our suggested definition of entity-level SA.

2.1 Related work on short texts

Mitchell et al. (2013) introduce the task of identifying which occurrences of named entities are sentiment targets in a text, and further to classify the sentiment towards these as positive or negative. They narrow the scope of named entities to volitional entities; i.e. organizations and persons. Zhang et al. (2015) follow up on this work and expand on the TSA task description.

The work of Mitchell et al. (2013) is highly related to our task in that they for each text identify the sentiment polarity towards each volitional entity mentioned in the text. However, while Mitchell et al. (2013) work on Twitter data, hence by their

nature very short texts, our goal in this work is to extend the task to longer texts where the sentiment towards one volitional entity may be expressed through multiple mentions and opinion expressions.

On a similar note, Jiang et al. (2011) also analyze sentiment towards each target in tweets. The goal of their work is, for a corpus of tweets and a query term, to return tweets classified as positive or negative towards the entity in the query term. The work goes beyond situations where the named entity is the sentiment target, into what is described as "extended targets". This work is a step in the direction of locating an extended set of segments of the text that are relevant for the sentiment towards each entity.

2.2 Related work on longer texts

One possible way of aggregating sentiment expressions is through the application of Coreference Resolution (CR) techniques. De Clercq and Hoste (2020) explore the benefits of CR in Aspect-Based Sentiment Analysis (ABSA). They find that when adding gold-standard coreference information, their models increase their performance by up to one percentage point, while using automatically retrieved CR data (Lee et al., 2013; De Clercq et al., 2011) would have no, or even negative effect on the performance of their ABSA models. Stoyanov and Cardie (2006) explore the benefits of holder coreference resolution as a method for extracting the document-level sentiment expressed by one holder. The task of target coreference resolution is mentioned in suggestions for future work. Farra and McKeown (2017) address the task of open-domain TSA, where their goal is to cluster targets and identify salient entities towards which opinions are expressed in the text. These targets may be nouns, noun phrases, events or concepts.

Steinberger et al. (2017) present their Europe Media Monitor (EMM) system, where the overall task is to detect the positive or negative sentiment towards persons and organizations. They process about 70 languages and therefore use linguistically light-weight methods that can work with low-resource languages.

The recent work of Luo et al. (2022) addresses some of the same limitations in previous work as observed by us, namely the limited focus on identifying sentiment targets only at the level of sentences or tweets. A new multi-domain dataset is

provided, annotated for sentiment targets and polarity, where all targets referring to the same entity are joined in a nested target structure. Their dataset provides a step towards targeted sentiment analysis at the document-level. There is no special treatment of volitional entities however.

2.3 Terminology review

Of the few papers we could find that use the term "entity-level sentiment analysis", several of them appear to treat it as synonymous to targeted SA (TSA) (Li and Lu, 2017; Alimova and Tutubalina, 2019; Huang and Fang, 2020; Sweeney and Padmanabhan, 2017; Engonopoulos et al., 2011). In TSA, the task is for each sentence to extract any segment being the target of a sentiment expression, and the sentiment polarity towards this target. This notion of a target is in line with the target expressions in widely used datasets for TSA (Pontiki et al., 2014, 2016), where there is no linking or aggregation to the document-level. Having the TSA term to describe this task, we suggest that the ELSA term is a better fit for an aggregated entity-level, where one entity may be linked to several targets in multiple sentences.

We did, however, also find a few studies employing the term "entity-level sentiment analysis" about entities aggregated from several mentions in the text (Farra and McKeown, 2017; Steinberger et al., 2017; Luo et al., 2022). These papers were presented in Section 2.1.

2.4 Conclusions from the literature study

We find that while there has been quite some work on classifying the polarity of particular occurrences of named entity mentions in text, like that of Mitchell et al. (2013) and Zhang et al. (2015), we have seen only a few studies that attempt to link or cluster several related segments of a document to resolve sentiment towards entities at the document-level. From existing research, we have seen that coreference resolution (Stoyanov and Cardie, 2006), PMI (Jiang et al., 2011) and semantic clustering (Farra and McKeown, 2017) have been used to resolve which entity mentions belong together. None of these papers were found reporting on actual experiments for the task of assigning one sentiment classification per volitional entity per document though. The recent paper by Luo et al. (2022) represents our closest match.

While we find that our usage of the term "entity-level sentiment analysis" is thematically related to

a few other usages in the literature, we do not see any established competing use of the term. We therefore suggest ELSA as an appropriate and descriptive term for the task discussed in the current paper.

3 The sentiment dataset

In order to investigate how the ELSA task relates to pre-existing sentiment analysis tasks, we wanted to build on an existing full-text document collection that is annotated for sentiment information at several levels of analysis. The suite of annotated datasets that are based on the Norwegian Review Corpus – NoReC (Velldal et al., 2018) – fits this bill. While Section 3.2 describes how we build on NoReC to create an exploratory dataset for ELSA, we first describe the different levels of existing annotations in NoReC below.

3.1 NoReC

NoReC is a multi-domain dataset of full-text professional reviews published in Norwegian online news sources, and a subset of the documents have been annotated for fine-grained and sentence-level sentiment. Each review in NoReC is accompanied by a rating given by the reviewer, on a scale from 1 to 6. We here take this to serve as a polarity label for the overall document.

Fine-grained sentiment A subset of NoReC has been annotated for fine-grained sentiment information in NoReC_{fine} (Øvrelid et al., 2020), including holders, target expressions, polar expressions, polarities, and polar intensities. This was one of the datasets used in the recent SemEval shared task on structured sentiment analysis – SemEval-2022 Task 10 (Barnes et al., 2022).¹ As presented in Table 1, the NoReC_{fine} training split consists of 327 documents, comprising 8634 sentences, giving an average of 26.4 sentences per document. The training split contains 5000 unique sentiment targets. Based on this data we can derive datasets for TSA as well as sentence-level SA, as described below.

Target-level sentiment We here describe how we derive a dataset for targeted SA – dubbed NoReC_{tsa} – on the basis of NoReC_{fine}. A given sentiment target in NoReC_{fine} may be the target of multiple opinion expressions, each with different polarity and intensity. We assign a value from 1 to

¹<https://competitions.codalab.org/competitions/33556>

3 to the sentiment intensities (‘slight’, ‘standard’ or ‘strong’), and assign the sum of all sentiments to each target. This sum is clipped to a scale from -3 (strong negative) to 3 (strong positive) for each target. The data is made available on GitHub.²

Sentence-level sentiment We also derive a 4-class sentence-level sentiment dataset from NoReC_{fine}. We take each sentence in NoReC_{fine} with positive opinions only to be a positive sentence (and vice versa for negatives). Sentences without any opinion expressions are labeled "Neutral", and sentiments with both positive and negative sentiments are labeled "Mixed".

Document-level sentiment The document-level sentiment polarity is derived from the rating provided by the review author. Table 1 shows the distribution of documents and sentences relative to these ratings for the subset of data included in the NoReC_{fine} training split and in our exploratory ELSA dataset (as further detailed below). For the same two respective subsets of data, Table 2 shows the distribution of documents and sentences across different domains of reviews.

As we can see, the dataset is highly unbalanced. The extreme ratings of 1 and 6 are rare, and are in our experiments therefore merged with their adjacent ratings. We arrive at sentiment polarity classifications for each document by labeling ratings 1 and 2 as negative, rating 5 and 6 as positive, and we here categorize ratings 3 and 4 as neutral polarity.

In sum, the ecosystem of annotations derived from NoReC_{fine} provides us with a multi-domain dataset of full-text reviews annotated for sentiment at multiple levels of analysis; document-level, sentence-level and target-level. This allows for comparing different strategies for aggregating SA information from different levels as estimators of entity-level sentiment. In order to evaluate these strategies and information sources, the next section describes how we create an exploratory ELSA dataset, adding information about entity-level sentiment for a subset of the documents in NoReC_{fine}.

3.2 An exploratory dataset for ELSA

To create an evaluation dataset for ELSA, we sample 50 documents at random from the 327 documents in the NoReC_{fine} train split. There are in total 1345 sentences in this evaluation set.

²https://github.com/lrgoslo/norec_tsa

Rating	NoReC _{fine} Train		ELSA subset	
	Docs	Sents	Docs	Sents
1	8	151	1	32
2	27	475	7	121
3	62	1345	4	98
4	91	2504	15	399
5	109	3225	19	586
6	30	934	4	109
Total	327	8634	50	1345

Table 1: Distribution of review ratings (1–6) in the full NoReC_{fine} training split and our ELSA subset.

Category	NoReC _{fine} Train		ELSA subset	
	Docs	Sents	Docs	Sents
games	16	445	2	62
literature	35	877	5	148
misc	1	36	1	36
music	111	1915	13	246
products	30	1753	6	298
restaurants	6	290	1	44
screen	118	2920	20	449
sports	2	149		
stage	8	249	2	62
Total	327	8634	50	1345

Table 2: Distribution of sentences and documents across the multiple domains in the NoReC review corpus.

As a pre-processing step, Named Entities are extracted from the texts using the Huggingface pipeline and the pretrained ScandiNER.³ language model, fine-tuned for NER on all Nordic languages, where the Norwegian part of the training data is provided by the NorNE corpus (Jørgensen et al., 2020). The reported scores for Norwegian NER are good, $\approx 91\%$ F₁. Since our goal is to identify sentiment towards volitional entities only, we keep only the entities with PER and ORG label. We find the NER model to perform well, with few or none missed entities. During manual inspection, spurious entities were deleted, and the few misclassifications were corrected.

Resolving coreference by substring matching

A volitional entity may have several entity mentions in the text. In order to cluster entity mentions in a text, we resolved the various mentions of a

³Available at <https://huggingface.co/saattrupdan/nbailab-base-ner-scandi>

I met John Wayne yesterday. We said hello on the street when he was taking his grandchild for a walk. John is such a nice guy. Nothing like Clint, who is still very handsome, but seems quite arrogant.	
John Wayne [John Wayne, John]	Positive
Clint [Clint]	Negative

Figure 1: Example of entity mentions in text as positive/negative sentiment targets, together with the overall entity-level sentiment labels.

volitional entity by simple substring matching, and also check for the genitive marker *-s* in Norwegian, such that both *Jo Nesbø* and *Nesbøs* are resolved to the same entity. One of the few errors introduced by this approach was that *Elisabeth I* was resolved to the same entity as *Elisabeth II*. These errors were subsequently corrected.

3.2.1 Manual ELSA labeling

For each unique entity in the text that the entity mentions represent, we manually evaluate the documents' sentiment towards these entities into the categories "Positive", "Negative", or "Neutral". For entities that are targets of both positive and negative sentiment expressions, we consider what the document as a whole conveys. This sentiment labeling was performed by one reader, after several readings, when necessary. Figure 1 shows a constructed example text, with the extracted entities, their entity mentions, and their manual sentiment classification.

We can conclude from reading the text that the sentiment towards *John Wayne* is positive, while the sentiment towards *Clint* is negative. There is one positive sentiment expression towards *Clint*: *handsome*, and one negative expression: *arrogant*. As readers of the entire text, we perceive the overall sentiment towards *Clint* to be negative. The entity *John Wayne* has two entity mentions in the text: *John Wayne* and *John*. The pronoun *he* is not an entity mention, but an anaphor, coreferential with *John Wayne*. The grandchild is not an entity because it is not named.

By manually assessing the sentiment expressed towards each entity in the 50 selected documents, we arrive at a dataset of volitional entities and the sentiment towards them at the document-level. The

	ORG	PER	#	%
Pos	8	76	84	30%
Neg	3	25	28	10%
Neu	36	131	167	60%
Total	47	232	279	100%

Table 3: The sampled subset of 50 documents from NoReC_{fine} contains the names of 279 volitional entities. The manual annotations assigned a neutral entity-level sentiment to the majority of these.

Mentions per entity	Entities	Entity mentions
1	188	188
2	39	78
3+	52	266
Total	279	532

Table 4: The evaluation data contain 279 volitional entities with an average of 1.9 entity mentions per entity. 33% of the entities have more than one mention.

majority of the entities are neutral, as shown in Table 3. This distribution is in line with Mitchell et al. (2013) who find that for their Spanish twitter dataset 24% of the entities are positive targets, 16% of the entities are negative targets, and 61% of the entities are neutral.

4 Analysis of target-independent approaches to ELSA

If all volitional entities mentioned in a positive text are the target of positive sentiment, the task would be limited to an overall sentiment classification for the text. This naïve approach serves as a baseline for further studies. In the following, we compare our manually labeled entity-level polarities with those of the documents they appear in, as well as those of each sentence with a corresponding entity mention.

4.1 ELSA polarity vs. document polarity

As mentioned in Section 3, each document is assigned the overall polarity positive, neutral or negative, based on its review rating.

In Figure 2 and Table 5 we present the results of an analysis of entity polarities for the different document ratings. We find that only 47.7% of the ELSA entities have the same polarity as the corresponding document-level label. We further observe




Figure 2: Relative distribution of entity polarities for the different document ratings. Neutral entities are in majority across all ratings.

Rating	Entity polarities			Entities
	Neg	Neu	Pos	
1-2	15	29	3	47
3	3	12	4	19
4	5	53	24	82
5-6	5	73	53	131
Total	28	167	84	279
True pos				133
Accuracy				0.477

Table 5: Distribution of entity polarities for each of the document ratings categories. As ratings of 1 and 6 are rather scarce, we merge them with their adjacent rating.

that the neutral entities are quite evenly distributed – between 55% and 65% across all document polarities. These results clearly suggest that simply inferring entity-level polarity from the document-level is insufficient.

4.2 ELSA polarity vs sentence polarity

We now turn to the sentence-level polarity labels presented in Section 3. Since the ELSA entities may have multiple mentions, they may appear in multiple sentences. We here aggregate the polarity towards an entity by considering it as positive if it is mentioned in more positive than negative sentences (and vice versa). The polarity is considered neutral if the entity appears in only neutral sentences. When entity mentions are equally present in pos-

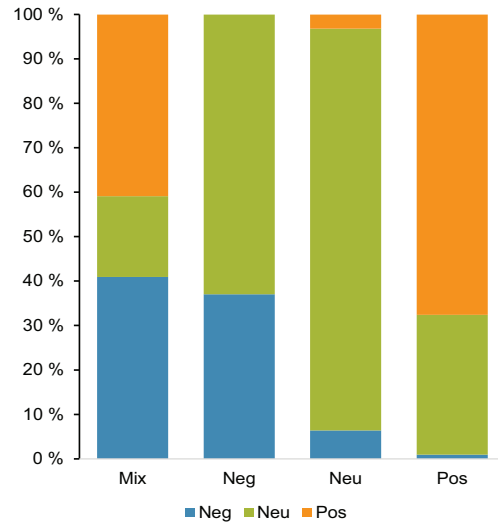


Figure 3: Relative distribution of entity polarities for each of the sentence sentiment categories.

Pol.	Entity polarities			Tot.
	Neg	Neu	Pos	
Mix	9	4	9	22
Neg	10	17		27
Neu	8	113	4	125
Pos	1	33	71	105
Total	28	167	84	279
True pos				194
Accuracy				0.695

Table 6: Distribution of entity polarities for each of the sentence sentiment categories. Approximately 70% of the entities were correctly labeled when using sentence polarity as a proxy.

itive and negative sentences, or in mixed polarity sentences only, mixed polarity is assigned.

The results are summarized in Figure 3 and Table 6. We find that neutral entities are more frequent than negative entities in sentences with negative polarity. An example of neutral entities that appear in a non-neutral sentence is provided in Example 1. The sentence is classified as negative and without any sentiment targets. The annotated sentiment towards *Julian Assange* is neutral.

- (1) Det gir en ganske merkelig effekt, litt som å treffe Julian Assange på Disneyland. *This has a quite peculiar effect, somewhat like meeting Julian Assange at Disneyland.*

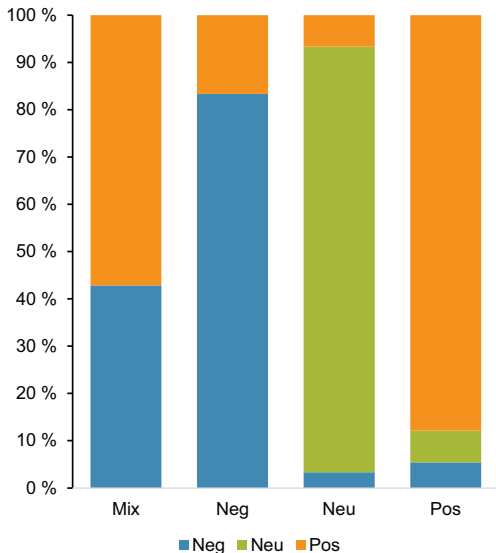


Figure 4: Relative distribution of entity polarities for each of the polarities derived from TSA.

116 of the entities (42%), have mentions in the same sentence as an entity of a conflicting polarity. The mixing of polarities inside a sentence indicates that inferring entity polarity from sentence polarity has limited potential. In sum, we find that 69.5% of the entities would be correctly resolved by inferring entity-level sentiment from the sentence sentiment.

5 A target-dependent approach to ELSA

An intuitively more promising approach is to derive ELSA sentiment labels from targeted sentiment analysis. Our TSA dataset presented in Section 3.1 is annotated for sentiment towards each target with a scale from -3 (strong negative) to 3 (strong positive). We aggregate these labels from the target-level to the entity-level for each entity by including only the sentiment targets that overlap with an entity mention, and summing the sentiment values for these targets. This approach leaves 7 entities with an unresolved classification due to equally strong positive and negative sentiment towards its entity mentions. These entities are placed in the "Mix" category.

As shown in Table 7, 229 of our 279 volitional entities, 82%, are given the correct sentiment label when aggregating the TSA annotations. This is an encouraging improvement over previous results. Virtually all neutral entities receive a neutral label

TSA	Entity polarities			Tot.
	Neg	Neu	Pos	
Mix	3		4	7
Neg	11		3	14
Neu	10	162	21	193
Pos	4	5	56	65
Total	28	167	84	279
True pos				229
Accuracy				0.821

Table 7: Distribution of entity polarities for each of the polarities derived from TSA.

Error type	#	%
Missing CR	2	0.72%
Missing TER	18	0.65%
Anno. disagreement	24	8.60%
Mixed polarity	6	2.15%

Table 8: Distribution of different error types when classifying entity sentiment based on aggregated TSA: Missing coreference resolution (CR), missing target–entity resolution (TER), annotator disagreement, and cases of ties from mixed TSA polarities.

through this approach. The 50 entities that are not correctly labeled using this approach, however, call for further analysis.

5.1 Studying the remaining 50 entities

In a further analysis step, we manually inspect the misclassified entities following TSA to look for common causes of errors.

Annotator disagreement Table 8 shows that human disagreement was the most important cause. For these instances, we have interpreted the sentences differently with respect to sentiment, than the original annotators of NoReC_{fine}. This serves as an example of the subjectivity and one may say fragility of human sentiment annotations. This is also evident in the moderate inter-annotator agreement for the NoReC_{fine} annotations, reported to be 73% F₁ for targets, when counting binary overlap (Øvrelid et al., 2020).

Coreference resolution To our surprise, only two (out of 50) entities were incorrectly classified due to lack of coreference resolution. One such case is presented in example (2) below.

- (2) Joseph Goebbels er på den måten i sjeldent, men ikke godt, selskap. Han er blitt stående som selve bildet på naziregimets hensynsløse og ondskapsfulle hatideologi. *Joseph Goebbels is in this respect in an exclusive, but not good company. He is the embodiment of the cruel and evil nazi regime's ideology of hatred.*

In Example 2, *Joseph Goebbels* in the first sentence is not annotated as a sentiment target in NoReC_{fine}, while *Han* 'He' in the second sentence is a negative sentiment target. Through coreference resolution, *Han* 'He' would likely be resolved to refer to *Joseph Goebbels*, and the sentiment towards the entity would be correctly resolved.

Sentiment-relevant relation extraction From our manual inspection we find another cause of misclassification pointing to the need for what we might call "target–entity resolution" rather than classical anaphora / cataphora coreference resolution. These are examples that require inference of semantic relations between different target expressions and entities in the text. Typical examples of this in our data are examples of a work of art where the sentiment towards the work of art implies a sentiment towards the creator, or a noun describing a group of people where an entity is a member.

- (3) Magnus Beite har skaffet filmen musikk som kler miljøet
Magnus Beite has provided the film with music that suits the environment
- (4) Bandet bestående av Daniel Birkeland på gitar [...] og trommis Helge Nyheim klarer virkelig å gjenskape Beatles [...] på en måte som det står respekt av.
The band consisting of Daniel Birkeland on guitar [...] and drummer Helge Nyheim really manage to give life to Beatles [...] in a way that commands respect.

In Example 3, *musikk* 'music' is a positive sentiment target. When reading the full text, we also perceive this as a positive sentiment towards *Magnus Beite*. In Example 4, *Bandet* 'The band' is the sentiment target, and we perceive this to imply positive sentiment towards the individual members of the band as well.

For these types of examples, if we had access to a "member of"-relation between the band mem-

bers and the band, and a "creator of"-relation between the photographer and the photograph, or more generally a relation of target–entity affiliation, the sentiment towards the volitional entities would have been resolved correctly. (Note that traditional aspect-categories would not be of help here, as we would still be missing the relations between the targets and the relevant entities.) These observations open new research questions, however, about which semantic relations are sentiment-relevant and under which circumstances. Jiang et al. (2011) also point to this question, with the example that a sentiment about someone's behavior usually means a sentiment about the person, while a sentiment about someone's colleague usually has nothing to do with the person.

6 Modeling

The above data analysis was performed using gold standard data for the various sentiment-levels (document, sentence and target). In order to gauge the performance attainable in a more realistic setting, we present results using automatically derived TSA information in the following.

Since we have no directly annotated ELSA training data, we create a baseline model using the proxy task of TSA. The exploratory ELSA dataset was taken from the training split of NoReC_{fine}, and we therefore join the remainder of the training split with the development split to create our baseline training data, and we perform the final evaluation on the 50 documents in our ELSA dataset. The model setup is adapted from the Huggingface example configuration for NER.⁴ We use default hyperparameters and perform no hyperparameter tuning. We therefore chose to not set aside data for a dev set, and not touch the original NoReC_{fine} test split, in order to allow for further research on this data split. We preprocess our training data the same way as we did with the ELSA data in Section 5, finding PER and ORG labels in the data through NER, and deriving sentiment towards these through the pre-existing TSA target annotations. Volitional entities that are sentiment targets, receive the sentiment label "Positive" or "Negative", while the volitional entities that are not sentiment targets are labeled "Neutral". With this setup, volitional entities and their sentiment polarity is predicted for

⁴<https://github.com/huggingface/transformers/tree/main/examples/pytorch/token-classification>

Pred entity pol	Gold entity pol			FP
	Neg	Neu	Pos	
Neg	4	0	4	1
Neu	13	154	27	35
Pos	9	6	48	8
Missed	2	7	5	

Table 9: A confusion matrix of gold and predicated entity polarities. The final column indicates false positive entity predictions. Our simple NER+TSA-based baseline model for ELSA predicted the name and sentiment polarity correctly for 206 of the 279 entities in the test data. The model produced 44 false positive entities, resulting in Precision: 66.7%, Recall: 73.8% and F_1 : 70.1%.

the 50 documents in the exploratory dataset. The predicted entity mentions were resolved document-wise through substring matching, before summing the predicted polarities. For evaluation, these predictions are compared with our manually resolved and annotated entities.⁵ Our evaluation ignores whether the model assigned the correct NER category PER or ORG. Our evaluation shows that this baseline model has a F_1 -score of 70.1%, as shown in Table 9.

7 Conclusion

In this paper we have explored the task of entity-level sentiment analysis (ELSA) – the task of determining the aggregated or overall (i.e. document-level) polarity expressed towards an entity in a text. In particular for longer text, which might comprise several distinct entities, entity mentions and opinions, this is a potentially complex task. The paper has surveyed existing literature and terminology in adjacent and related previous work, in addition to presenting an exploratory data analysis to shed more light on what is required to solve the task.

Relating to the latter point, in order to assess the relevance to ELSA of existing approaches to sentiment analysis at different granularities – i.e. document-, sentence-, and target-level SA – we perform a task analysis on the basis of a Norwegian multi-domain review dataset containing all these layers of sentiment annotation. When utilizing the fine-grained sentiment annotations of the

⁵Code for our data collection, analysis and modeling will be available at <https://github.com/egilron/elsa-introduction>

NoReC_{fine} dataset, we found that the overlap between a volitional entity and TSA annotations, gave us the correct sentiment category for 82% of our 279 manually evaluated ELSA entities. We further found that, using automated TSA, we obtained an F_1 -score of 70.1%. TSA therefore appears to be highly relevant, though not sufficient, for the ELSA task.

For the ELSA classifications that could not be derived from NER and TSA annotations, we found surprisingly few cases that would be correctly resolved through coreference resolution. This is in line with the findings of De Clercq and Hoste (2020). Our dataset is likely too small, however, to draw any definitive conclusions about the importance of coreference resolution for ELSA. On the other hand, we did observe a need for more generally resolving relations that tie target expressions to their corresponding entities.

In our exploratory data analysis, we have identified at least four categories of information that could potentially benefit the classification of a document’s overall sentiment towards entities: (i) named entity recognition, (ii) targeted sentiment analysis, (iii) coreference resolution, and (iv) what we dub target–entity resolution. The latter being a concept derived from working with the examples in our dataset, and refers to the task of identifying which entity a given target expression relates to. This appears to be a missing link in a pipeline for ELSA based on TSA, and to explore methods for filling this gap is a suggestion for further work.

Acknowledgements

The work documented in this publication has been carried out within the NorWAI Centre for Research-based Innovation, funded by the Research Council of Norway (RCN), with grant number 309834, and the SANT project (Sentiment Analysis for Norwegian Text), also funded by the RCN (grant number 270908).

References

- Ilseyar Alimova and Elena Tutubalina. 2019. [Entity-level classification of adverse drug reactions: a comparison of neural network models](#). In *Proceedings of the 2019 Workshop on Widening NLP*, pages 132–134, Florence, Italy. Association for Computational Linguistics.
- Jeremy Barnes, Laura Ana Maria Oberländer, Enrica Troiano, Andrey Kutuzov, Jan Buchmann, Ro-

- drigo Agerri, Lilja Øvrelid, and Erik Velldal. 2022. SemEval-2022 task 10: Structured sentiment analysis. In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*, Seattle. Association for Computational Linguistics.
- Orphee De Clercq and Veronique Hoste. 2020. [It's absolutely divine! can fine-grained sentiment analysis benefit from coreference resolution?](#) In *Proceedings of the Third Workshop on Computational Models of Reference, Anaphora and Coreference*, pages 11–21, Barcelona, Spain (online). Association for Computational Linguistics.
- Orphée De Clercq, Véronique Hoste, and Iris Hendrickx. 2011. [Cross-domain Dutch coreference resolution](#). In *Proceedings of the International Conference Recent Advances in Natural Language Processing 2011*, pages 186–193, Hissar, Bulgaria. Association for Computational Linguistics.
- Nikos Engonopoulos, Angeliki Lazaridou, Georgios Paliouras, and Konstantinos Chandrinou. 2011. ELS: a word-level method for entity-level sentiment analysis. In *Proceedings of the International Conference on Web Intelligence, Mining and Semantics (WIMS '11)*, pages 1–9, Sogndal, Norway.
- Noura Farra and Kathy McKeown. 2017. [SMARTies: Sentiment models for Arabic target entities](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 1002–1013, Valencia, Spain. Association for Computational Linguistics.
- Zhihong Huang and Zhijian Fang. 2020. An entity-level sentiment analysis of financial text based on pre-trained language model. In *2020 IEEE 18th International Conference on Industrial Informatics (INDIN)*, volume 1, pages 391–396. IEEE.
- Long Jiang, Mo Yu, Ming Zhou, Xiaohua Liu, and Tiejun Zhao. 2011. [Target-dependent Twitter sentiment classification](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 151–160, Portland, Oregon, USA. Association for Computational Linguistics.
- Fredrik Jørgensen, Tobias Aasmoe, Anne-Stine Ruud Husevåg, Lilja Øvrelid, and Erik Velldal. 2020. [NorNE: Annotating named entities for Norwegian](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 4547–4556, Marseille, France. European Language Resources Association.
- Heeyoung Lee, Angel Chang, Yves Peirsman, Nathanael Chambers, Mihai Surdeanu, and Dan Jurafsky. 2013. Deterministic coreference resolution based on entity-centric, precision-ranked rules. *Computational linguistics*, 39(4):885–916.
- Hao Li and Wei Lu. 2017. Learning latent sentiment scopes for entity-level sentiment analysis. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, San Francisco, USA.
- Yun Luo, Hongjie Cai, Linyi Yang, Yanxia Qin, Rui Xia, and Yue Zhang. 2022. [Challenges for open-domain targeted sentiment analysis](#).
- Margaret Mitchell, Jacqui Aguilar, Theresa Wilson, and Benjamin Van Durme. 2013. [Open domain targeted sentiment](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1643–1654, Seattle, Washington, USA. Association for Computational Linguistics.
- Lilja Øvrelid, Petter Mæhlum, Jeremy Barnes, and Erik Velldal. 2020. [A fine-grained sentiment dataset for Norwegian](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 5025–5033, Marseille, France. European Language Resources Association.
- Maria Pontiki, Dimitris Galanis, Haris Papageorgiou, Ion Androutsopoulos, Suresh Manandhar, Mohammad AL-Smadi, Mahmoud Al-Ayyoub, Yanyan Zhao, Bing Qin, Orphée De Clercq, Véronique Hoste, Marianna Apidianaki, Xavier Tannier, Natalia Loukachevitch, Evgeniy Kotelnikov, Nuria Bel, Salud María Jiménez-Zafra, and Gülşen Eryiğit. 2016. [SemEval-2016 task 5: Aspect based sentiment analysis](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 19–30, San Diego, California. Association for Computational Linguistics.
- Maria Pontiki, Dimitris Galanis, John Pavlopoulos, Haris Papageorgiou, Ion Androutsopoulos, and Suresh Manandhar. 2014. [SemEval-2014 task 4: Aspect based sentiment analysis](#). In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pages 27–35, Dublin, Ireland. Association for Computational Linguistics.
- Ralf Steinberger, Stefanie Hegele, Hristo Tanev, and Leonida Della Rocca. 2017. [Large-scale news entity sentiment analysis](#). In *Proceedings of the International Conference Recent Advances in Natural Language Processing, RANLP 2017*, pages 707–715, Varna, Bulgaria. INCOMA Ltd.
- Veselin Stoyanov and Claire Cardie. 2006. [Partially supervised coreference resolution for opinion summarization through structured rule learning](#). In *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing*, pages 336–344, Sydney, Australia. Association for Computational Linguistics.
- Colm Sweeney and Deepak Padmanabhan. 2017. [Multi-entity sentiment analysis using entity-level feature extraction and word embeddings approach](#). In *Proceedings of the International Conference Recent Advances in Natural Language Processing, RANLP 2017*, pages 733–740, Varna, Bulgaria. INCOMA Ltd.
- Erik Velldal, Lilja Øvrelid, Eivind Alexander Bergem, Cathrine Stadsnes, Samia Touileb, and Fredrik Jørgensen. 2018. [NoReC: The Norwegian review corpus](#). In *Proceedings of the Eleventh International*

Conference on Language Resources and Evaluation (LREC 2018), Miyazaki, Japan. European Language Resources Association (ELRA).

Meishan Zhang, Yue Zhang, and Duy-Tin Vo. 2015. [Neural networks for open domain targeted sentiment](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 612–621, Lisbon, Portugal. Association for Computational Linguistics.