

## ERC Advanced Grant 2019

### Part B2: The scientific proposal

## Subjective Language: Its means and ends

### (SUBLIME)

Subjective language is the language we use to communicate our opinions, emotions and evaluations: about good and bad, right and wrong, beautiful and ugly, and countless variations on these themes; with words whose semantic values depend on *our* values, not just on states of affairs but also on states of mind. As a first illustration, contrast the statement (1) with the statement (2):

- (1) This was the largest audience to ever witness an inauguration ...
- (2) Kennedy's short sentences and simple words make his inaugural the best in history.

The size of an inauguration audience is a matter of objective fact. It may be difficult to measure it, but all will agree that it is possible in principle. The quality of an inaugural address, however, is a matter of opinion, insofar as two persons can disagree over it and neither be wrong.

That languages are not only tools for exchanging information but also for exchanging opinion, with dedicated resources for either task, *viz.*, objective and subjective language, may seem a truism. However, in semantics, a systematic bias towards objective language has left a serious gap in our understanding of what subjective language is and how it works. The SUBLIME project sets out to redress this imbalance, decisively advancing our grasp of subjectivity in language and our ability to recognize its forms and determine its functions.

This goal is challenging but nevertheless achievable due to recent developments in semantics, partly pioneered by the PI, and because the PI and his team will be able to combine methods and models which have not yet been brought together to bear on subjectivity. The outcomes will be a more sensitive semantics and pragmatics, a more solid basis for public discourse about language, and, owing to the many faces and interfaces of subjectivity, new perspectives for text and discourse analysis, sociolinguistics, and language technology.

### Section a. State-of-the-art and objectives

Subjective language is a stepchild in the study of natural language meaning. Semantic theory and its associated methodology are geared towards objective language, and the insight that subjective language plays a distinctive role has been late in coming; it is, in fact, only some fifteen years old.

The main reason for this is that formal semantics, the dominant theory of linguistic meaning, owes its basic tools to logic, where sentence meaning is standardly modeled with truth conditions. For objective meaning, this makes good sense: there is a truth of the matter, and this truth can in principle be established by inspecting the world, or possible worlds which are accessible from it. For example, the truth value of (1) above could be determined by counting 45 crowds since 1789.

Subjective meaning is a less good fit: to the extent that it makes sense to talk of true or false, telling one from the other is more of a discretionary matter, depending not only on the world but also on a subject's – usually the speaker's – outlook on it. Thus (2) above cannot be evaluated by measuring sentences and classifying words – its value depends on a subjective evaluation.

Discussion of how best to model this dependency began in earnest in the early 2000s, with the work by Kölbl (2004), MacFarlane (2005) and Lasersohn (2005). To be sure, a distinction had long been made between objective and subjective content, in particular regarding moral judgments, from Thomas Aquinas (*Summa Theologiae*) to Gibbard (1990). But integrating subjective content into the paradigm of formal semantics is a project of the present millennium, and it has generated a rich line of research with many innovations, like the theory formulated by Coppock (2018).

Still, this line of research is hampered by three severe limitations.

First, there is far from a consensus on how best to model the special dependency that subjective content is evidently subject to; nor is there a consensus on what expressions have subjective content or on whether the same dependency is in play in every case. Since Lasersohn (2005) formulated a ‘relativist’ semantics where words like *tasty* depend on a judge supplementing the time and world indices of evaluation, many scholars have proposed stating the same idea in different terms (Bouchard 2012, Bylinina 2017, among others), or argued for substituting a ‘contextualist’ theory (Pearson 2013, Silk 2016, among numerous others). Other proposals defend so-called ‘expressivism’ (Huvenes 2014) or define subjectivity as a pragmatic phenomenon (Kennedy and Willer 2016, *i.a.*). Some of the mentioned theorists take a narrow view of what expressions carry subjective meaning, concentrating on taste predicates; some take a wider view, whereas yet others, like Kennedy (2013), distinguish between narrow and wide concepts of subjectivity, or, like Vardomskaya (2018), argue that subjectivity is a context sensitive phenomenon in which diverse parts of language can share.

The state of the semantic study of subjectivity is thus rather immature, and it is a leading idea in the project that this is largely due to the second and third of the three mentioned limitations.

The second limitation concerns the empirical base of the field: the inventory of ways to express subjective content has not been properly explored. Most studies have focused on select adjectives, while in fact, it is clear that nouns, as in (3a), and verbs, as in (3b), are parts of the picture.

- (3) a. Quite honestly, your opinions on Brexit are a nonsense.  
 b. Scotland has deserved the right to hold a second referendum.

Besides, once we look beyond formal semantics to less theoretically oriented neighbor disciplines – see below – various bodies of work suggest that the scope of subjectivity resources is even wider, spanning more levels of a language than its lexicon. Evaluations and opinions are evidently not only conveyed through words but also through constructions, discourse markers, intonation, and more. Connectedly, the diagnostics that can be used to tell whether an expression has subjective content have not yet been investigated on the scale called for in pioneering work by Sæbø (2009).

The third limitation on the study of subjectivity in formal semantics is its isolation from other lines of research which, while more descriptively oriented or more oriented towards applications, are essentially about the same object: Subjective language. In the following fields, subjective language is a focus of interest primarily in virtue of the *functions* it can have:

- Corpus linguistics and functional linguistics (see, e.g., Hunston and Sinclair 2000),
- critical discourse analysis and sociolinguistics (see, e.g., Englebretson 2007, White 2015),
- computational linguistics, specifically ‘sentiment analysis’ (see, e.g., Deng and Wiebe 2016).

These strands of research have passed ‘under the radar’ of the formal semantic study of subjectivity, and the other ways around. In consequence, knowledge about the pragmatics of subjective language, concerning the uses it is put to, is fragmented and less comprehensive and precise than it could be.

Overall, existing research leaves three overarching lacunae of knowledge:

- ◇ The lacuna of the *means* of subjective language: we do not know the range of its sources, or how to distinguish subjectivity from objectivity or among different kinds of subjectivity.
- ◇ The lacuna of the *essence* of subjective language: as yet, no consensus has formed on what a good theory of the semantics and pragmatics of subjective language looks like.
- ◇ The lacuna of the *ends* of subjective language: we lack accurate knowledge about how speakers and writers use it to convey their values and to shape those of hearers and readers.

Filling these lacunae will bring clear gains to the theory of natural language meaning, as well as to applied branches of linguistics such as discourse analysis, sociolinguistics, and language technology; moreover, it will lay groundwork for a better informed public discourse about linguistic practices.

The three knowledge lacunae provide the project with structure: it will have one research area on the ‘means’ side and another on the ‘ends’ side, mediated by a third, theoretical area at the centre. These three areas will, in turn, be subdivided into three to four sub-areas, as depicted in Figure 1.

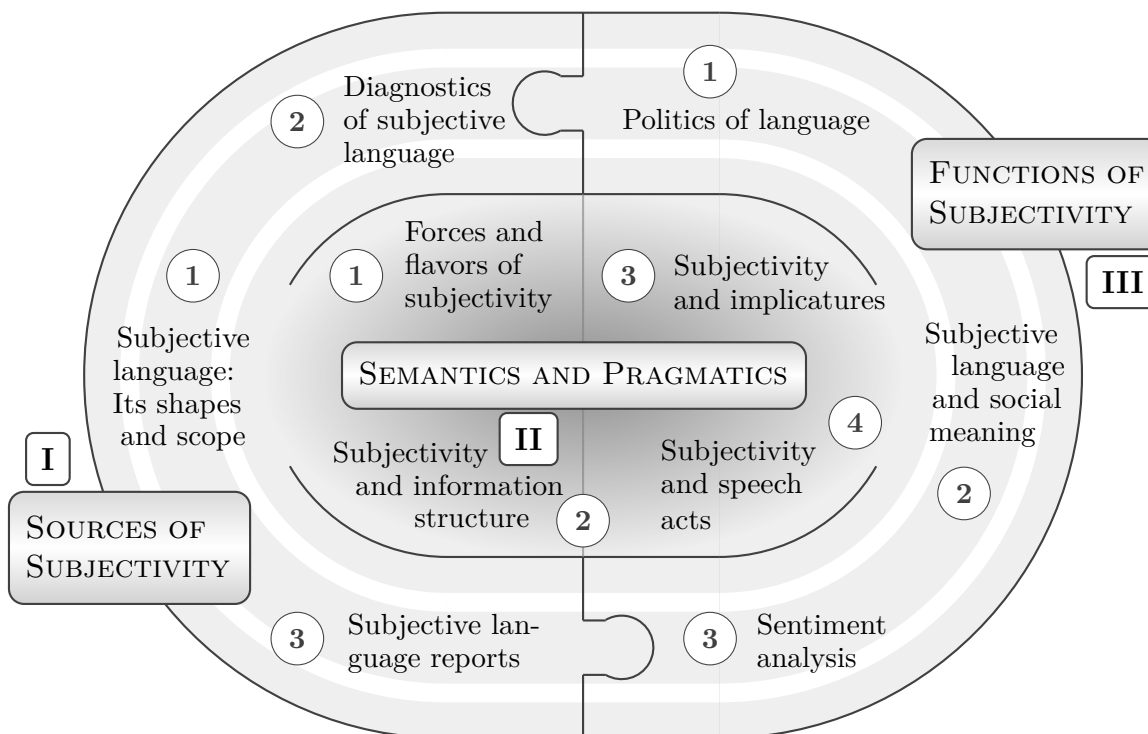


Figure 1: The areas and sub-areas of research and their interconnections

With reference to the Lasswellian communication model – ‘Who says what, how, to whom, with what effect’ (Lasswell et al. 1952), area I is concerned with the *how* and the form side of the *what*, area II is concerned with the content side of *what* and with the intentional side of *with what effect*, and area III is concerned with the *who*, the *to whom* and, in particular, with the *with what effect*.

In area I, the overarching aim is to chart the landscape of the ‘built-in emoticons of language’ and to show that across languages, a variety of previously understudied indicators can help identify a wide and varied family of sources for subjective content, as illustrated in Figure 2.

Area II is committed to theory-building on the basis of novel evidence flowing from area I and to reappraisals of received models of at-issue/non-at-issue content, conversational implicatures, and illocutionary acts – speech acts – in the light of subjective locutions and illocutions.

In area III, the searchlight is directed onto the perlocutions of subjective language to uncover, on the basis of the theoretical advances under way in area II, what works how in political discourse, what senders communicate what attitudes to what receivers, and what lessons are to be learned across the boundaries between sentiment analysis and the other sub-areas of the project.

Once these desiderata are realized, the study of linguistic meaning will be in an improved state, both in itself, in relation to neighboring fields and in terms of public outreach and broader impact. On the one hand, the theory will be more balanced, with less of a bias towards objective language, and new and more favorable perspectives will open for the relevant branches of applied linguistics. On the other hand, and as importantly, the theory will be better equipped to inform public debate about what language reveals and how language forms attitudes, because these issues evolve around language which is predominantly subjective, where linguistics has traditionally had little to offer.

The three areas and ten sub-areas of research will now be described in more detail.

## I SOURCES OF SUBJECTIVITY

The prime aim of the work to be done in this area is to establish

- what expressions or other linguistic resources are subjective,
- in what ways these different linguistic resources are subjective,
- by what criteria they are subjective in the way they are subjective.

There are large gaps in our knowledge on these issues, for although we know more about the shapes that subjective language takes today than twenty years ago, even scholars who, like Umbach (t.a.), look “beyond *fun* and *tasty*” tend to focus on lexical resources, predominantly English adjectives. It may well be that wide swaths of the landscape of subjective language remain to be charted.

Even the tools with which this can be done are underexplored. To tell subjective and objective content apart, ‘judge arguments’ and ‘subjective attitude verbs’ (Sæbø 2009) are used as criteria, augmenting the less tangible notion of ‘faultless disagreement’ (see page 6 on area I, sub-area 2). But again, scholars have tended to restrict attention to a narrow range of English exponents.

Even so, it has been noted, in particular by Kennedy (2016), that such criteria may differ and that different criteria can give different results. This raises the prospect that there is not only one way in which a linguistic resource can be subjective, and that such variation is correlated with the variation in subjective linguistic resources. In fact, there is preliminary evidence that

- subjective content can be expressed with a wide and varied family of linguistic resources, at different levels, from the ‘micro-level’ of the lexicon to the ‘macro-level’ of discourse,
- different resources can be subjective in different ways; in particular, there is reason to differentiate between evaluative and non-evaluative subjectivity,
- this can be demonstrated through a wider range of indicators than hitherto appreciated.

The overarching research objective in area I is to show that this is indeed the case, and how.

### I ① *Subjective language: Its shapes and scope*

According to Alba-Juez and Thompson (2014), subjectivity – or, in their own term, *evaluation*; the expression of sentiments and stances in discourse – has many ‘faces’: It permeates language across levels and can be expressed, explicitly and implicitly, through a wide range of linguistic resources.

Some of these faces are very familiar – for example, adjectives like *good*, *bad* and *ugly* – and generally, words in the major parts of speech are ‘usual suspects’ for subjective content, although examples of nouns or verbs are much more rarely cited; (4) and (5) are from Zillig (1982: 137):

- (4) Müllers Diskussionsbeitrag war Quatsch. (German)  
Müller’s discussioncontribution was nonsense
- (5) Müller hat gefaselt. (German)  
Müller has rambled

The resources at this ‘micro-level’ are items that can form the core predication in a sentence.

One step up, at a ‘meso-level’ of whole sentences or items that can apply at sentence level, we encounter a variety of modal expressions, like the modal adverb *troppo* ‘too’ as it occurs in (6):

- (6) Questo film è troppo centrato sull’amore. (Italian)  
this film is too centered on-the-love

This is also the locus of some intensifiers (Beltrama 2018) and of evaluative adverbs, which express a stance towards the core proposition (Bonami and Godard 2008, Mayol and Castroviejo 2013), and it is where optatives and exclamatives, analyzed by Grosz (2012) in terms of an evaluative operator, and many so-called Mad Magazine sentences (Akmajian 1984) are at home.

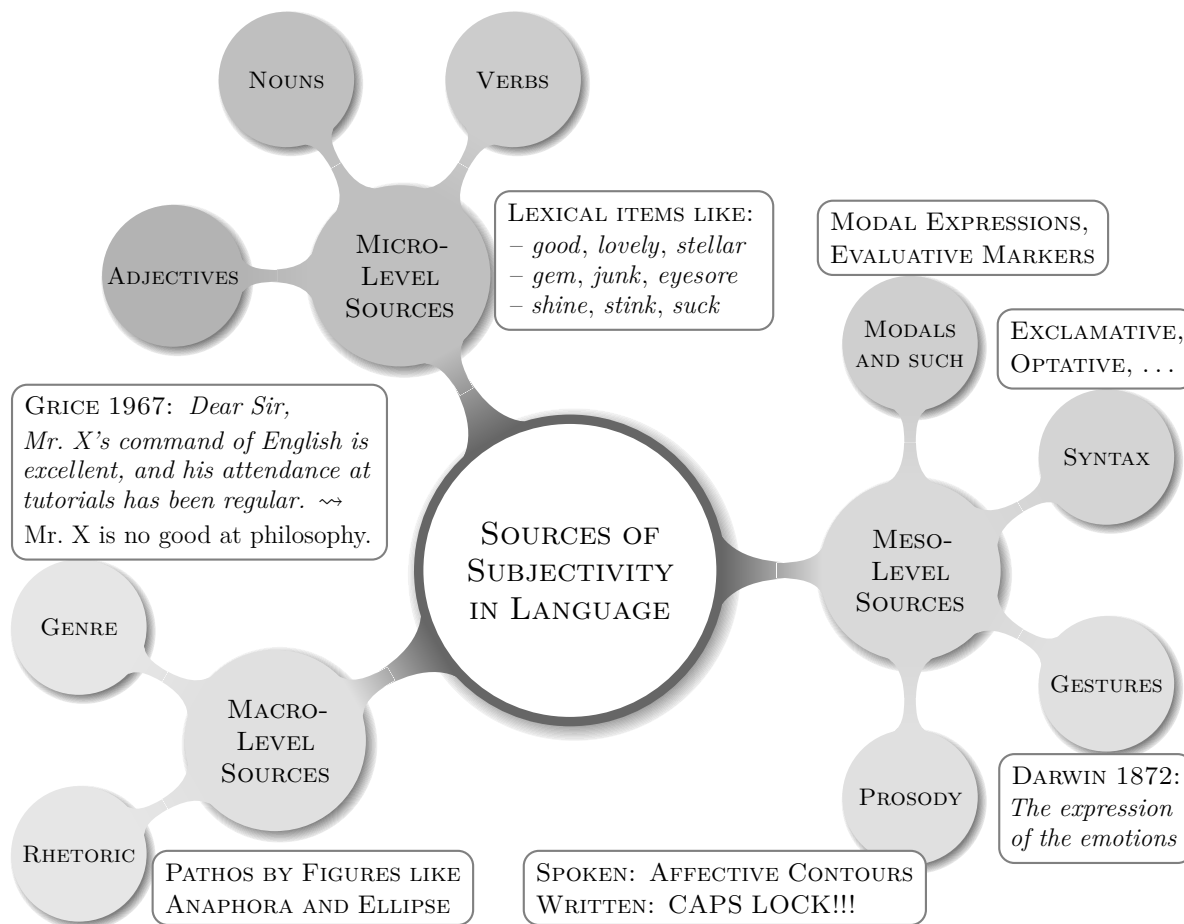


Figure 2: A mindscape of subjectivity resources

In this diagram of the subjectivity sources that are to be researched, some nodes are a lighter gray than others; this is to reflect that those areas are less well charted than the darker ones.

Moving now from the near-top right, still relatively dark gray area of exclamative and optative clauses downward to the suprasegmental areas of gestures and prosody, we enter just such terrain: underexplored means of expressing emotions, evaluations and opinions. The time may be right: There is a current surge in research on ‘co-speech’ and ‘pro-speech’ gestures generally (see, e.g., Schlenker 2019), and particularly facial gestures would seem to play a central role for evaluativity. As for prosody, recent developments (see, e.g., Wagner 2016) suggest that intonation encodes an ‘emotional valence’, and in particular, that pitch, stress and rhythm can convey evaluative content even when the lexical items and the other segmental resources are in themselves objective.

At the level of the larger discourse, finally, subjective meaning can evidently also be conveyed through the ways we cut the cake of information and wrap it – through the text’s build and type. First, as taught by the Sophists, rhetorical figures like ἀναφορά *anaphora* and ἔλλειψις *ellipsis* are figures of pathos, subtle means of adding an evaluative dimension to an otherwise neutral content. (7) is a case in point: in context – *Official apology to women who had relationships with German soldiers during the Second World War* by Norway’s Prime Minister Erna Solberg, October 23, 2018 – the fragmentary style implicitly sends the message that an official apology is overdue.

- (7) It divided sisters and brothers. Parents and daughters. Whole families were branded.  
For generations.

Second, a range of text genres – letters of recommendation, obituaries, promotional text, reviews – are in themselves evaluative, coming with expectations of evaluative language, something which can ‘rub off’ on the interpretation of basically non-evaluative statements, positively or negatively.



## I ② Subjective language: Diagnostics

Telling subjective from objective language would be a matter of subjective judgment if it were not for certain tests to help us decide. The most well-known such test is that of *faultless disagreement* – if you say that wolves are cute and I say no they are not, we disagree but neither of us is at fault. However, using this to probe for subjectivity is itself to some degree a discretionary matter; indeed, [Solt \(2018\)](#) reports experimental evidence that judgments whether a disagreement is faultless or not vary by degrees across (non-)evaluative adjectives. Instead, what will be under scrutiny here are two more tangible but underexplored diagnostics: *Judge phrases* and *subjective attitude verbs*.

The former appear as prepositional phrases, as in (8), or as dative phrases, as in (9).

(8) That’s for me the most beautiful tune ever written. (Joe Henderson, about *Lush Life*)

(9) Der Film ist i h r z u sentimental, ... (German; *dativus iudicantis*)  
the film is she.DAT too sentimental, ...

Questions about judge phrases in need of answers include what kinds of subjectivity they diagnose – thus [Bylinina \(2017\)](#) argues that *мне* in (10) is not a judge argument but an experiencer argument.

(10) М не э то т фильм был неинтересен. (Russian)  
I.DAT this film was not.interesting

Subjective attitude verbs are propositional attitude verbs that require subjective propositions; thus if they are felicitous, that is a sign that the complement clause contains something subjective. Clear cases are the verbs *synes* in Danish and Norwegian and *tycka* in Swedish, the Norwegian Sign Language verb shown in Figure 4, and the Chinese verb *jué dé*.



Figure 3: Subjective attitude verb in Chinese

The main source on subjective attitude verbs is [Sæbø \(2009\)](#); early sources include [Ducrot \(1975\)](#). See [van Wijnbergen-Huitink \(2016\)](#) for a survey.




Figure 4: Subjective attitude verb in NSL

In English, there is evidently a division of labor between three verbs: *find* (which in the relevant sense can only embed small clauses), *feel* and *think* (evidently ambiguous between an objective sense and a subjective sense).

(11) is infelicitous due to the objective content of the complement clause:

(11) #Growing numbers of Americans find the Earth flat.

What this sensitivity consists in is a contentious issue; [Hirvonen \(2014: 52ff.\)](#) rejects that data from *find* can be used to test for subjectivity. But in any case, comparisons of attitude verbs intra- and interlinguistically are likely to bring out telling nuances and facets of subjective meaning.

## I ③ Subjective language: Reports

There is another, more indirect way in which subjective language is reflected inside language itself, namely with speech act reports. There are designated verbs for representing subjective talk or text: For me to be correctly described as *criticizing* or *praising* someone for something, or as *hailing* or *railing* (against) somebody or something, I must use subjective language, or at any rate use language in a subjective sense. In (12), subjective language, in quotes, forms part of the report:

(12) Kremlin slams “witchhunt” election hacking allegations...

This phenomenon is understudied (see, however, [Fabricius-Hansen and Sæbø 2011](#)). Studying it – again, across a variety of languages – is likely to provide another key to results on the shades and flavors of subjective language, results which can inform the work to be done in area II.

## II THE SEMANTICS AND PRAGMATICS OF SUBJECTIVE LANGUAGE

Empirical work only takes us so far: Filling white spaces on the map of the sources of subjective language, uncovering descriptive generalizations, drawing distinctions and establishing correlations are all good and necessary steps, but leave a void in the explanatory dimension. Here as elsewhere, the theorist’s task is to sharpen our vague notions and fit our findings into well-founded frameworks where they make sense. Pre-theoretic concepts like ‘evaluation’ and ‘opinion’ are mere labels until we model them in formal semantic terms, and subjectivity itself is an elusive category as long as we lack a clear idea of the appropriate mode of interpretation for subjective language. Also, it is plausible that subjective language exists for a purpose, or for multiple purposes, but to understand what those are, we need to engage with pragmatic theory.

In turn, standard versions of semantic and pragmatic theory may be due for radical reappraisals once subjective language is taken seriously.

### II ① *Forces and flavors of subjectivity*

Scholars concur that the semantic values of subjective parts of language depend on something else or more than what the world is like. But how best to model this dependence is a pending question, ‘relativist’ positions vying with ‘contextualist’ positions as well as with each other and yet others; this list offers a cross-section of recent relevant work: [Coppock 2018](#), [Gutzmann 2016](#), [Kennedy and Willer 2016](#), [Lasersohn 2016](#), [Silk 2016](#). We will seek to inform this question by introducing two new distinctions: subjective **force** and subjective **flavor**.

Subjective force is most easily illustrated in a relativist framework. Assume that – in principle – the semantic value of any expression is fixed by a world of evaluation and by a ‘judge’ of evaluation, and, for convenience, that worlds and judges form finite sets. Figure 5 shows a toy evaluation space where the rows are worlds and the columns are judges and where (i) an **objective** proposition, (ii) a **weakly subjective** proposition, and (iii) a **strongly subjective** proposition is gray:




Figure 5: A toy evaluation space with horizontal worlds and vertical judges

A strongly subjective proposition is both true and false in all of the worlds, depending on the judge, a weakly subjective proposition is true and it is false in some of the worlds, depending on the judge; an objective proposition is one for which judges do not make a difference in any world.

Subjective flavor is the dimension that is to encode the content of terms like ‘emotion’, ‘opinion’ and ‘evaluation’ and relate them to categories like ethics, aesthetics and taste. – Not accidentally, the terms ‘force’ and ‘flavor’ are borrowed from the theory of modality ([Kratzer 1991, 2012](#)): there is a strong affinity between subjectivity and modality which has not yet been given due attention.

### II ② *Subjectivity and information structure*

Objective and subjective content sometimes cohabitate: One utterance contains portions of both. Preliminary work ([Sæbø 2017](#)) shows that this need not result in a picture like (ii) (see Figure 5); if the objective material is *not at-issue*, as often reflected in the utterance’s information structure ([Beaver et al. 2017](#)), the picture will be like (iii) for those worlds where the proposition is defined. Thus (13) merely presupposes and does not entail that Maret Sara has written a book.

- (13) Su mielas lea Maret Sara čállán #(miellagiddevaš) girjji. (North Saami)  
 (s)he.GEN mind.LOC is Maret Sara written interesting book.ACC  
 ‘She finds the book Maret Sara has written interesting.’

Conversely, non-restrictive subjective modification fails to count toward a proposition’s subjectivity. Sæbø (2017) argues that the role of information structure in partitioning content into objective and subjective can supply criteria for choosing among theories – and this will be a key stepping-stone.

## II ③ *Subjective language and implicatures*

Conversational implicatures, the core of (Neo-)Gricean pragmatics, are mostly studied with reference to objective language, where the notion of informativeness is singularly apt. Scalar implicatures are relevant for subjective language too (“What was it like?” “It was okay.”  $\rightsquigarrow$  it wasn’t great), but in addition, what can be implicated with subjective means goes beyond what can be implicated with objective means. Therefore, the study of subjectivity can serve to fill out the picture of implicature, and the theory can help us better understand what things can be done with subjective language.


The idiom “veiled message” is a rather accurate catchphrase for the implicatures at issue here: Saying or writing something neutral can, in the right circumstances, convey something value-laden, and saying or writing something positive can actually convey something negative, and vice versa. Irony (see, e.g., Grice 1967: Lecture II, 17; Garmendia 2018) is one relevant phenomenon.

## II ④ *Subjective language and speech acts*

Speech Act theory holds that every utterance has a purpose and the purposes form natural classes. But, from Austin (1962) via Searle (1969), Sadock (1974), Searle (1975a), and numerous others to Krifka (2017) and the ongoing ERC Advanced Grant project *Speech Acts in Grammar and Discourse*, no consensus has formed on the classification of illocutionary acts. One reason is that one utterance may show characteristics of two types of acts – and the predominant view since Searle (1975b) has been that in such cases, one is **primary** while the other is **secondary**.

This conception of speech acts corresponds to the model on the left in Figure 6: Any utterance has a dominant illocutionary force, belonging in one and only one class.

Subjective language may pose a challenge to this model. In view of utterances typically made with (14) or (15), it is not easy to maintain clear distinctions between assertives, expressives and directives (see Roberts (2018) for a recent reappraisal of the Searlian taxonomy):

- (14) That’s a very nice drawing!  
 (15) Isn’t my granddaughter just so precious 

One goal of the work to be done in this sub-area is to build evidence for the alternative model on the right in Figure 6, with seamless transitions between the archetypes and room for hybrid acts.



Figure 6: Discrete (left) and dense (right) topology of speech acts



### III FUNCTIONS OF SUBJECTIVITY

Speech acts are generally taken to play out at three levels: (1) the *locution*, what is said or written; (2) the *illocution*, the intended effect of the locution; and (3) the *perlocution*, the locution’s actual effect. Conventionally, the focus of attention has been on the two first levels, the utterance itself and the intention behind it. But in particular with regard to subjective language, there are strong reasons to direct the searchlight onto the perlocutionary consequences of talk and text.

Three sets of issues merit attention here. First, locutions, and especially subjective locutions, often do not take place one by one and in isolation from each other but in a complex public space, where voices and countervoices influence people’s minds in ways which may or may not differ from the *il*-locutions. There is a need to know more about *what works how*.

Second, effects of subjective talk and text do not occur in social vacua but will vary with the *who* and the *whom* of Lasswell’s classical formulaic model of communication (Lasswell et al. 1952). More should and can be known about how social meaning manifests itself in subjectivity and about what senders communicate what attitudes to what receivers.

Third, the vibrant field of Sentiment analysis in Language Technology testifies to the fact that bodies of subjective text mean much to bodies of individuals; both directly, influencing them, and indirectly, in that they hold valuable keys to knowing our attitudes and predicting our behavior.

In short, there is a gap between the formal and the functional, the theoretical and the practical, which this third section of the project aims at bridging.

#### III ① *Subjectivity and the politics of language*

As argued by Beaver and Stanley (t.a.), when linguists make the standard idealization that hearers and speakers, readers and writers follow the Gricean Cooperative Principle and use language toward a common goal, they exclude from consideration ways of using language that are not strictly communicative but nevertheless highly consequential. A case in point is the way that metaphors are often used, consciously or unconsciously, as a means of dehumanization: When members of a group are described in terms that are selectionally restricted to lower organisms or natural forces, it serves to define that group as an out-group and to reduce empathy for its members. Or, when opponents are characterized in negative terms, it will color the perception of them within the in-group:

- (16) Csak jön egy olyan agyament ótvar hazug idióta, Ulrike Lunacek, ... (Hungarian)  
 just come a such lobotomized scab liar idiot, Ulrike Lunacek, ...  
 ‘Along comes this imbecile, impetiginous lying idiot Ulrike Lunacek, ...’

Evaluative elements of language are crucial tools in such practices. Hence, when Beaver and Stanley urge the need to refocus the theory of meaning on political speech and on **speech practices**, then a designated focus on subjective content, its sources and potencies ought to further their program. This defines the main objective here: Contribute to improved knowledge of how we are affected by political discourse by zooming in on evaluative elements and uncovering what works how.

The importance of this line of work is amplified by claims that public language use is changing, from a more objective to a more subjective, oral style. Thus Beckett (2017) discusses the ‘affective turn’ in journalism and the “increasingly subjective, relativist identity political discourse”, as well as “the increasing role of emotion in the creation, dissemination and reception of journalism in general and political communications in particular”. This trend is linked to new media technologies, which have “enabled and amplified the shift to subjectivity”; see also, e.g., Papacharissi (2015).

These issues are largely debated in media studies and critical discourse analysis, and the notions ‘affective’, ‘subjective’, *et cetera* are not grounded in formal definitions or in testable criteria. This motivates a more specific objective in regard to subjectivity and the politics of language: Contribute to the debates about a subjective turn in public discourse by bringing in a linguistic perspective so as to provide objective criteria for assessing the claims that are made.

### III ② *Subjective language and social meaning*

That language is not just a formal system but also a social medium is something of a truism. But there is a specific sense in which language has social significance: It shows who we are, or would like to be, and who we perceive our discourse partners to be. Choices we make when speaking or writing can carry a *social meaning* (see, e.g., [Burnett 2017](#) and [Eckert 2012](#)).

There are signs that this can be especially pronounced when subjective language is used, and in particular, that a bond is forged between subjective and social meaning by the concept of *stance* (see, e.g., [Englebretson 2007 \(ed.\)](#) and [Jaffe 2009 \(ed.\)](#)), the way that speakers position themselves in terms of evaluation or social relation. For instance, different ways to express aesthetic assessment may signal that we assume or defer authority, or social pressure may sway us to express insincere assessments. H. C. Andersen's 1837 tale of *the emperor's new clothes* provides an extreme example: Everybody but the child introduced towards the end professes to adore the nonexistent attire.

- (17) Ja er det ikke magnifique! (Danish)  
yes is it not magnificent

From slightly different perspectives, [Beltrama \(2018\)](#) and [Wiltschko, Denis and d'Arcy \(2018\)](#) have argued that certain discourse markers (*totally* in American English and *eh?* in Canadian English) with clear affinities to subjectivity are at the same time socially meaningful.

As interesting as the social meaning/subjectivity intersection appears in light of the cited work, the question whether there is indeed a special bond between subjective and social content, and if there is, what this bond consists in, has not been systematically examined. To do that is the prime task set for the research in this sub-area. A more specific sub-goal is to study the well-worn notion 'faultless disagreement' (see p. 5) under the lens of stance-taking to show that the degree to which discussants treat disagreements over subjective judgments as faultless is sensitive to social relations.

### III ③ *Subjectivity and sentiment analysis*

For about as long as subjective language has been a topic in semantics and philosophy of language – 15–20 years – it has been a subject of interest in computational linguistics and language technology, and Sentiment Analysis, or Opinion Mining, remains a vibrant field of research and development.

The reason is that there are needs for systems that can detect the stances we take in texts we type and sort the texts accordingly, by recognizing clues of subjectivity; a few key sources are [Liu 2010](#), [Pang and Lee 2008](#), [Wiebe et al. 2004](#). This is done with increasingly sophisticated methods (see [Barnes, Klinger and Schulte im Walde 2017](#), [Bratlie and Bakken 2016](#), [Deng and Wiebe 2016](#) for some select examples). Yet, the research is generally not conducted in a context that integrates semantics and pragmatics (see [Benamara, Taboada and Mathieu 2017](#) and [Taboada 2016](#)), though there is reason to assume that an integrated perspective will be fruitful, in particular in regard to precision. As a case in point, distinguishing between expressions of opinion and reports of opinion – say, between (18a) and (18b) – is a challenging task for a system.

- (18) a. Censeo Carthaginem esse delendam. (Latin)  
consider.1P.SG Carthage.ACC be destroy.GER.ACC  
'It is my opinion that Carthage should be destroyed.'
- b. Cato censet Carthaginem esse delendam.  
'It is Cato's opinion that Carthage should be destroyed.'

The work to be done here will consist in clearing paths between sentiment analysis and the other (sub-)areas in the project (see [West, Paskov, Leskovec and Potts \(2014\)](#) on sentiment analysis and social meaning) to show that there are mutual lessons to be learned and, in particular, how results from area I can be utilized for the benefit of more accurate systems.

## Section b. Methodology

Four approaches are suitable for achieving the research objectives identified in the preceding section. These are, enumerated and supplied with color codes for ease of reference:

- 1** Fact finding based on corpus exploration
- 2** Fact finding based on experiment and elicitation
- 3** Theory building based on found facts
- 4** Cross-pollination of theories

The first and second are essential for providing empirical input to the project from outside, while the third and fourth utilize the interconnections between the ten sub-areas to ensure a flow of fact and theory inside it. The former two methodologies will be central in more descriptive sub-areas, whereas more theoretical sub-areas will primarily be researched with the latter two; see Figure 7.

### **1** Fact finding based on corpus exploration

Text corpora have been a well-tapped source for linguistic fact-finding for some thirty years, and even elementary search techniques remain useful. But the last two to three decades have also seen the evolution of increasingly sophisticated tools for extracting linguistic information from corpora, some of which bring central research objectives within closer reach, particularly in area I.

Concretely, to identify the sources of subjectivity in a language, it will be advantageous to take a distributional semantic perspective (see [Boleda and Herbelot 2016](#)) and extract the distribution patterns of subjectivity diagnostics with the analysis tool [Sketch Engine](#) (see [Kilgarriff et al. 2014](#)); as different diagnostics give rise to divergent patterns, this method will help to differentiate among the various sources and thus to draw a fine-scale map of the landscape of subjective language.

There are many such landscapes, though, and to obtain a picture of what varies or is constant cross-linguistically, a selection of similar and dissimilar languages familiar to members of the team will be targeted for investigation and comparison. The plan is to focus on Mainland Scandinavian, German and two to four non-Germanic languages – one Romance and one Slavonic, see page 13 – and to employ translation mining methodology (see [van der Klis, Le Bruyn and de Swart 2017](#)) to create so-called semantic maps from multilingual, parallel corpora.

### **2** Fact finding based on elicitation and experiment

Some objectives will best be served with facts found through language users' reactions to stimuli. Thus the corpus-based study of select languages in area I will be supplemented by a broader survey deployed through the linguist-sourced database [Terraling](#). More often, though, elicitation methods are going to be suitable for answering research questions in areas II and III.

Some experimental work on subjective meaning exists: [Kaiser and Lee \(2017, 2018\)](#), [Solt \(2018\)](#), and [Verheyen, Dewil and Égré \(2018\)](#) collect judgments on different aspects of adjectives that are assumed to be more or less subjective, in the form of responses to two- or multiple-choice questions.

These studies use non-real-time techniques rather than real-time ones like eye-tracking or neuro-imaging, and that will be the case in this project as well; on the other hand, pure processing tasks where only the input is linguistic will be complemented by tasks where the output is linguistic too. Tasks combining processing and production are especially apt when the effects, actual or intended, of using subjective language are on the agenda, to elicit what these effects are perceived to be.

This approach is reminiscent of the methodology of Empirical Semantics pursued by Arne Næss in the 1930s (see, e.g., [Næss 1938](#) and [Chapman 2011](#)), using questionnaires and interviews to learn what we mean by key terms. Such results are less easy to quantify than fixed-choice responses, but, as shown by Næss, they can also be made amenable to quantitative, statistical analysis.

### 3 Theory building based on found facts

The theory-building work in area II will be building on a flow of evidence from fact-finding work. In addition, neighboring disciplines where relevant fact-finding has been done – corpus linguistics and computational linguistics, *inter alia* – will be providing valuable materials. Consequently, a much wider array of facts will be taken into account than in any comparable effort so far.

Good theory building requires not only plentiful materials but also effective tools and compliance with strict standards. The work will hold to the ideals upheld in formal semantics and pragmatics – formal rigor, mindfulness of counterevidence, compositionality, parsimony and explanatory power. Tools are to be found in the box of the formal semantic and pragmatic literature, first and foremost, in the form of proposals in recent work cited on page 7 and 8, and in the hypotheses outlined there about the proper ways to model subjective content and the uses subjective content can be put to.

### 4 Cross-pollination of theories

Subjective language has been studied from different theoretical angles: the formal linguistic angle, the functional linguistic angle, and the angle of discourse analysis, without many points of contact. Efforts to mutually inform theories across the boundaries of these sub-disciplines of linguistic study will constitute a methodology of its own.

Specifically, the formal study of subjective meaning will stand to gain from integrating insights from (i) Stance theory in the framework of critical discourse analysis and sociolinguistics (e.g., [Du Bois 2007](#)), (ii) Appraisal theory in the functional linguistics paradigm (e.g., [White 2015](#)), and (iii), in the context of political language, Framing theory (see, e.g., [Lakoff and Wehling 2016](#)).

Conversely, perspectives from formal semantics and pragmatics are likely to lead to a sharpened focus on the key concepts in those theories and to a clearer understanding of their significance.

### 5 Some illustrations

A few concrete illustrations of how the methods described above will be deployed are in order.

- In area I, sub-area 2, the primary objective is to track and compare diagnostics of subjectivity across a selection of languages and thereby bring out differences among types of subjectivity; this objective will be pursued through a combination of three modes of corpus exploration: (i) probing for diagnostics by tracing co-occurrence patterns of subjective items, (ii) probing for natural classes of subjective items by tracing co-occurrence patterns of diagnostics, (iii) tracing (dis)similarities in the inventories of diagnostics and subjective items in multilingual corpora. Pioneering work by [Lande \(2009\)](#) has shown that this is a fruitful approach.
- In area II, sub-area 4, the central goal is to build a case that utterances made with sentences with subjective items as active ingredients often do not fit into one single speech act type but motivate a model where speech act types blend into each other. This goal will be approached with experimental tools in the spirit of [Beltrama \(2018a\)](#), eliciting judgments about responses, complemented by formats where not just judgments are elicited but also authentic responses and where subjects formulate answers to questions about the perceived point(s) of utterances made with declarative, exclamative or interrogative subjective-content sentences in contexts.
- In area III, sub-area 2, the key objective is to test the claim that subjective language is more likely than objective language to carry social meaning. Area III is generally where the greatest risk lies, because the actual functions of language are new ground for formal semantics and pragmatics; therefore, it is essential to bring together various theories and methods for cross-pollination. Concretely, the plan is to analyze and utilize the notions of [Stance](#) and [Appraisal](#) and, taking cues from [Drager \(2018\)](#), to adapt experimental procedures from sociolinguistics to designs which specifically elicit perceptions of social meaning along scales of subjectivity.

## 6 Participants and plan

The tasks that consist in reaching the objectives set out in Section **a** with the methods described above will be performed by a team where the members have clearly defined, yet interlocking roles. Beside the principal investigator, who will be devoting 60% of his time to the project throughout its 5-year duration, six scholars will be engaged full-time for periods ranging from 3 to  $4\frac{1}{2}$  years: One researcher ( $4\frac{1}{2}$  years), two postdocs (4 years and 3 years) and three PhD students (3 years each). The postdocs and the PhD students, mentored and supervised by the PI, will be working within and across sub-areas which abut and sometimes intersect with each other, and the researcher to be engaged for the project's full duration will also be sharing responsibility for its overall progress.

Somewhat more specifically, the three PhDs will be assigned with individual projects in area I (PhD1: *Non-lexical expression of subjective content*, PhD2: *Subjectivity diagnostics across languages*) and in area III (PhD3: *Sentiment analysis in light of semantic theory*); one postdoc (postdoc 1) will carry responsibility for the pragmatics of subjective language (area II), and another (postdoc 2) will be addressing issues concerning its political and social dimensions (area III). The senior researcher will, beside engaging with issues in areas I and II, act as the project's day-to-day coordinator, and the PI will, alongside his project management responsibilities, devote his efforts to collaboration, mentorship and supervision in and across all the ten sub-areas of the project, with an emphasis on semantics and pragmatics as far as collaborations are concerned.

The senior researcher position is intended to be filled by Alexandra Spalek, who earned her doctorate at the UPF, Barcelona, in 2014 and has since held a postdoc at the University of Oslo. Spalek is ideally suited for this role because (i) she is a specialist in lexical semantics, (ii) she has extensive experience in corpus linguistic methodology, (iii) she has strong teambuilding skills, and (iv) in addition to her native language German, she is fully fluent in Polish, Spanish and Catalan. With the PI's native competence in Norwegian and full fluency in German, core team members will thus be well poised to investigate subjectivity in and across a fair sample of closely and less closely related languages. For the postdoc and PhD positions, diverse and outstanding candidates will be recruited through international searches where key criteria include familiarity with corpus and/or experimental methodology or with text mining in Natural Language Processing.

The work to be done in the project faces two special challenges. Firstly, it involves an unusual variety of methods and objectives along the spectrum between basic and applied linguistic research. Secondly, the bulk of it will be carried out by staff who are hired for the purpose. Both challenges place high demands on coordination and mentorship. The PI's dedication to and experience with promoting early careers will be essential for meeting these demands, as will the plan to entrust a senior researcher with a general management role. Also, the team will not work in a vacuum but in a large and fertile environment of scholarship and research in linguistics, computational linguistics, and philosophy of language, across four departments. Notably, the PhD project *Sentiment analysis in light of semantic theory* will benefit from contact with collaborators in the project *Sentiment Analysis for Norwegian Text* (SANT, 2018–21), based at the Department of Informatics.

Strengthening ties with select senior scholars at the forefront of relevant research will further enhance the project's feasibility. The leading figures listed below share a keen interest in its agenda and bring together a unique blend of expertise, and these contacts will be promoted through guest researcher stays in Oslo extending from one week to two months.

- **Regine Eckardt**, Universität Konstanz, Germany – a leading figure in semantics, pragmatics and the semantics/pragmatics interface, and an ideal second supervisor for PhD1,
- **Louise McNally**, Universitat Pompeu Fabra, Barcelona, Spain – a leading figure in semantics and an expert in corpus method; an ideal second supervisor for PhD2,
- **Stephanie Solt**, Leibniz-Zentrum Allgemeine Sprachwissenschaft (ZAS), Berlin, Germany – a successful semanticist highly competent in experimental methodology.



Figure 7 lays out the schedule of the project. The work in the ten sub-areas, ordered vertically, is plotted along the 20 quarters of the 5-year runtime horizontally, in line with the four methodologies represented by their color codes. Each colored bar thus stands for a clearly defined research task; one colored grid square corresponds to appr.  $\frac{1}{4}$  person-year.




Figure 7: The sub-areas of research laid out as time-phased tasks with milestones

The numbered diamonds represent milestones – key intermediate objectives and control points in the form of (i) deliverables; reports on which to base abstracts for submission to key conferences or manuscripts for submission to journals, such submissions themselves, or materials for dissemination to wider audiences, (ii) PhD midway assessments and completed dissertations, and (iii) workshops where the progress of the project and the quality and quantity of its output can be assessed.

- ① Workshop on progress in area I · submissions to conferences and journals · PhD1 and PhD2 underway assessments
- ② Submission of journal article authored by researcher and PhD1 · dissertation submission 1 (Non-lexical expression of subjective content)
- ③ Submission of journal article authored by PI and PhD2 · dissertation submission 2 (Diagnostics of subjectivity across languages)
- ④ Submission on reports of subjective speech to conference
- ⑤ Conference Subjective content and information structure
- ⑥ Submission of journal articles authored by PI, researcher, PhD1, PhD2
- ⑦ Workshop The pragmatics of subjective language · submission to conference (postdoc 1)
- ⑧ Submission of journal article authored by PI and postdoc 1
- ⑨ Workshop The socio-politics of subjective content · submissions to broad-audience journals
- ⑩ Conference Subjective language in the public domain · submission of journal article authored by PI and postdoc 2 · joint book manuscript areas I–III
- ⑪ PhD3 midway assessment · workshop on non-heuristic strategies for sentiment analysis
- ⑫ Dissertation submission 3 (Sentiment analysis in light of semantic theory)



Supplementing these milestones, a denser mesh of joint activities and checkpoints will further help structure and integrate the different tasks within and across the three research areas:

- Team workshops, each with one of the three associated senior scholars attending, devoted to each of the three methodologies 1, 2 and 4 and how they are best to be deployed in the tasks,
- semi-annual retreats where team members build the mutual trust that allows them to benefit optimally from each other’s experience and expertise, and work out the mid-term agenda,
- weekly project meetings where the team reviews and adjusts the short-term agenda, with special regard to publications, including but not limited to those defined in the milestones.

## References

- Akmajian, Adrian (1984) Sentence types and the form-function fit. *Natural Language and Linguistic Theory* 2(1), 1–23.
- Alba-Juez, Laura and Geoff Thompson (2014) The many faces and phases of evaluation. In Geoff Thompson and Laura Alba-Juez (eds.), *Evaluation in Context* (= Pragmatics and Beyond New Series 242). Amsterdam and Philadelphia: John Benjamins, 3–23.
- Austin, J.L. (1962) *How to do things with words*. Cambridge, Mass.: Harvard University Press.
- Barnes, Jeremy, Roman Klinger and Sabine Schulte im Walde (2017) Assessing state-of-the-art sentiment models on state-of-the-art sentiment datasets. In *Proceedings of the 8th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, Copenhagen: ACL, 2–12.
- Beaver, David, Craige Roberts, Mandy Simons and Judith Tonhauser (2017) Questions Under Discussion: Where information structure meets projective content, *Annual Review of Linguistics* 3, 265–284.
- Beaver, David and Jason Stanley (forthcoming) *Hustle: The Politics of Language*. Forthcoming at Princeton University Press.
- Beaver, David and Jason Stanley (t.a.) Towards a non-ideal philosophy of language. To appear in *New School Graduate Philosophy Journal*, special issue on Ordinary Language Philosophy.
- Beckett, Charles (2017) Towards an emotionally networked journalism. *Polis: Journalism and society at the London School of Economics* (blogs.lse.ac.uk/polis/2017/08/31/).
- Beltrama, Andrea (2018) Totally between discourse and subjectivity: Exploring the pragmatic side of intensification. *Journal of Semantics* 35(2), 219–261.
- Beltrama, Andrea (2018a) Subjective assertions are weak: exploring the illocutionary profile of perspective-dependent predicates. In Uli Sauerland and Stephanie Solt (eds.), *Proceedings of Sinn und Bedeutung 22, volume 1*, Berlin: ZAS, 161–173.
- Benamara, Farah, Maite Taboada and Yannick Mathieu (2017) Evaluative Language Beyond Bags of Words: Linguistic Insights and Computational Applications. *Computational Linguistics* 43(1), 201–264.
- Boleda, Gemma and Herbelot, Aurélie (2016) Formal Distributional Semantics: Introduction to the Special Issue. *Computational Linguistics* 42(4), 619–635.
- Bonami, Olivier and Danièle Godard (2008) Lexical Semantics and Pragmatics of Evaluative Adverbs. In Louise McNally and Chris Kennedy (eds.), *Adverbs and Adjectives: Syntax, Semantics, and Discourse*, Oxford: Oxford University Press, 274–304.
- Bouchard, David-Étienne (2012) *Long-Distance Degree Quantification and the Grammar of Subjectivity*. PhD thesis, McGill University.

- Bratlie, Terje and Patrik Bakken (2016) *Sentiment analysis in Norwegian political news – employing replicated methods and experimental features to understand a complex domain*. Master thesis, Norwegian University of Science and Technology.
- Burnett, Heather (2017) Sociolinguistic interaction and identity construction: The view from game-theoretic pragmatics. *Journal of Sociolinguistics* 21(2), 238–271.
- Bylinina, Lisa (2017) Judge-dependence in degree constructions. *Journal of Semantics* 34(2), 291–331.
- Chapman, Siobhan (2011) Arne Naess and Empirical Semantics. *Inquiry* 54(1), 18–30.
- Coppock, Elizabeth (2018) Outlook-based semantics. *Linguistics and Philosophy* 41(2), 125–164.
- Darwin, Charles (1872) *The Expression of the Emotions in Man and Animals*. London: J. Murray.
- Deng, Lingjia and Janyce Wiebe (2016) How can NLP tasks mutually benefit Sentiment Analysis? A holistic approach to Sentiment Analysis. In *Proceedings of the 7th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*.
- Drager, Katie (2018) *Experimental Research Methods in Sociolinguistics*. London: Bloomsbury.
- Du Bois, John (2007) The stance triangle. In Robert Englebretson (ed.), *Stancetaking in Discourse: Subjectivity, evaluation, interaction*, Amsterdam: John Benjamins, 137–182.
- Ducrot, Oswald (1975) Je trouve que, *Semantikos* 1, 63–88.
- Eckardt, Regine (2014) *The Semantics of Free Indirect Discourse*. Leiden: Brill.
- Eckert, Penny (2012) Three waves of variation study: The emergence of meaning in the study of sociolinguistic variation. *Annual review of Anthropology* 41, 87–100.
- Englebretson, Robert (ed.) (2007) *Stancetaking in Discourse: Subjectivity, evaluation, interaction*. Amsterdam: John Benjamins.
- Fabricius-Hansen, Cathrine and Sæbø, Kjell Johan (2011) Behabitive Reports. In Elke Brendel, Jörg Meibauer and Markus Steinbach (eds.), *Understanding Quotation*, Berlin: Mouton de Gruyter, 85–106.
- Garmendia, Joana (2018) *Irony*. London: Cambridge University Press.
- Gibbard, Allan (1990) *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Cambridge, Mass.: Harvard University Press.
- Grice, Paul (1967) *Logic and Conversation*. Lecture notes from William James Lectures at Harvard. Mimeograph, University of California, Berkeley.
- Grosz, Patrick Georg (2012) *On the Grammar of Optative Constructions*. Amsterdam: John Benjamins.
- Gutzmann, Daniel (2016) If expressivism is fun, go for it! Towards an expressive account of predicates of personal taste. In Janneke van Wijnbergen-Huitink and Cécile Meier (eds.), *Subjective Meaning. Alternatives to Relativism*. Berlin: de Gruyter. 21–46.
- Hirvonen, Sanna (2014) *Predicates of Personal Taste and Perspective Dependence*. PhD thesis, University College London.
- Hunston, Susan and John Sinclair (2000) A local grammar of evaluation. In Susan Hunston and Geoffrey Thompson (eds.), *Evaluation in Text*, Oxford: Oxford University Press, 74–101.
- Huvenes, Torfinn (2014) Disagreement without error. *Erkenntnis* 79, 143–154.
- Jaffe, Alexandra (ed.) (2009) *Stance: Sociolinguistic Perspectives*. Oxford: Oxford University Press.
- Kaiser, Elsi and Jamie Herron Lee (2017) Experience matters: A psycholinguistic investigation of predicates of personal taste. In Dan Burgdorf, Jacob Collard, Sireemas Maspong, and Brynhildur Stefánsdóttir (eds.), *Proceedings of Semantics and Linguistic Theory* 27, 323–339.

- Kaiser, Elsi and Jamie Herron Lee (2018) Predicates of personal taste and multidimensional adjectives: An experimental investigation. In Wm.G. Bennett et al. (eds.), *Proceedings of the 35th West Coast Conference on Formal Linguistics*, 224–231. Somerville, MA: Cascadilla Proceedings Project.
- Kennedy, Christopher (2013) Two sources of subjectivity: Qualitative assessment and dimensional uncertainty. *Inquiry* 56(2-3), 258–277.
- Kennedy, Christopher (2016) Two kinds of subjectivity. In Janneke van Wijnbergen-Huitink and Cécile Meier (eds.), *Subjective Meaning: Alternatives to Relativism*, Berlin: de Gruyter Mouton, 105–126.
- Kennedy, Chris and Malte Willer (2016) Subjective attitudes and counterstance contingency. In Mary Moroney, Carol-Rose Little, Jacob Collard and Dan Burgdorf (eds.), *Proceedings of SALT 26*, LSA: Open Journal Systems, 913–933.
- Kilgarriff, Adam, Vit Baisa, Jan Bušta, Miloš Jakubíček, Vojtěch Kovář, Jan Michelfeit, Pavel Rychlý and Vit Suchomel (2014) The Sketch Engine: ten years on. *Lexicography* 1(1), 7–36.
- Kölbel, Max (2004) Faultless Disagreement. *Proceedings of the Aristotelian Society* 104, 53–73.
- Kratzer, Angelika (1991) Modality. In Arnim von Stechow and Dieter Wunderlich (eds.), *Semantics. An International Handbook of Contemporary Research*, Berlin: de Gruyter, 639–650.
- Kratzer, Angelika (2012) *Modals and Conditionals: New and Revised Perspectives*. Oxford: Oxford University Press.
- Krifka, Manfred (2017) Epistemic, evidential and discourse modalities in commitment space semantics. Presentation at the XPrag workshop *Non-at-issue Meaning and Information Structure*, University of Oslo.
- Lande, Inna (2009) *Subjektive Einstellungsberichte im Deutschen: Einstellungsverben und beurteilungsabhängige Prädikate*. MA thesis, University of Oslo.
- Lakoff, George and Elisabeth Wehling (2016) *Your Brain's Politics. How the Science of Mind Explains the Political Divide*. Exeter: Imprint Academic.
- Lasersohn, Peter (2005) Context dependence, disagreement, and predicates of personal taste. *Linguistics and Philosophy* 28, 643–686.
- Lasersohn, Peter (2016) *Subjectivity and Perspective in Truth-Theoretic Semantics*. Oxford: Oxford University Press.
- Lasswell, Harold Dwight, Daniel Lerner and Ithiel de Sola Poole (1952) *The comparative study of symbols: An introduction*. Stanford, CA: Stanford University Press.
- Liu, Bing (2010) Sentiment analysis and subjectivity. In Nitin Indurkha and Fred Damerau (eds.), *Handbook of Natural Language Processing, Second Edition*, Boca: Taylor and Francis Group, 627–666.
- MacFarlane, John (2005) Making Sense of Relative Truth. *Proceedings of the Aristotelian Society* 105, 321–39.
- Mayol, Laia and Elena Castroviejo (2013) (Non)integrated evaluative adverbs in questions: A cross-Romance study. *Language* 89(2), 195–230.
- McNally, Louise and Isidora Stojanovic (2017) Aesthetic adjectives. In James O. Young (ed.), *Semantics of Aesthetic Judgements*. Oxford: Oxford University Press. 17–37.
- Næss, Arne (1938) Common-sense and truth. *Theoria* 4(1), 39–58.
- Pang, Bo and Lillian Lee (2008) Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval* 2(1-2), 1–135.
- Papacharissi, Zizi (2015) *Affective Publics: Sentiment, Technology, and Politics*. New York: Oxford University Press.

- Pearson, Hazel (2013) A judge-free semantics for predicates of personal taste. *Journal of Semantics* 30(1), 103–154.
- Roberts, Craige (2018) Speech acts in discourse context. In Daniel Fogal, Daniel W. Harris, and Matt Moss (eds.), *New Work on Speech Acts*, Oxford: Oxford University Press.
- Sadock, Jerrold (1974) *Toward a Theory of Linguistic Speech Acts*. New York: Academic Press.
- Sæbø, Kjell Johan (2009) Judgment ascriptions. *Linguistics and Philosophy* 32, 327–352.
- Sæbø, Kjell Johan (2017) Subjective content at sole issue. Presentation at the XPrag workshop *Non-at-issue Meaning and Information Structure*, University of Oslo.
- Schlenker, Philippe (2019) Gestural semantics. *Natural Language and Linguistic Theory*.
- Searle, John (1969) *Speech Acts*. London: Cambridge University Press.
- Searle, John (1975a) A taxonomy of illocutionary acts. In K. Gunderson (ed.), *Language, Mind and Knowledge*, Minneapolis, MN: University of Minnesota Press, 344–369.
- Searle, John (1975b) Indirect Speech Acts. In Peter Cole and Jerry L. Morgan (eds.), *Syntax and Semantics 3: Speech Acts*, New York: Academic Press, 59–82.
- Silk, Alex (2016) *Discourse Contextualism*. Oxford: Oxford University Press.
- Solt, Stephanie (2018) Multidimensionality, subjectivity and scales: experimental evidence. In Elena Castroviejo, Louise McNally and Galit Weidman Sassoon (eds.), *The Semantics of Gradability, Vagueness, and Scale Structure*. Language, Cognition, and Mind, vol 4. Cham: Springer, 59–91.
- Taboada, Maite (2016) Sentiment analysis: An overview from linguistics. *Annual Review of Linguistics* 2, 325–347.
- Umbach, Carla (t.a.) Evaluative predicates beyond *fun* and *tasty*. To appear in Lisa Matthewson, Cécile Meier, Hotze Rullmann and Ede Zimmermann (eds.), *Wiley’s Companion to Semantics*.
- van der Klis, Martijn, Bert Le Bruyn and Henriëtte de Swart (2017) Mapping the PERFECT via Translation Mining. *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2*, 497–502.
- van Wijnbergen-Huitink, Janneke (2016) Subjective meaning: An introduction. In Janneke van Wijnbergen-Huitink and Cécile Meier (eds.), *Subjective Meaning: Alternatives to Relativism*, Berlin: de Gruyter Mouton, 1–19.
- Vardomskaia, Tamara N. (2018) *Sources of Subjectivity*. PhD thesis, University of Chicago.
- Verheyen, Steven, Sabrina Dewil and Paul Égré (2018) Subjectivity in gradable adjectives: The case of tall and heavy. *Mind and Language*.
- Wagner, Michael (2016) How to be kind with prosody. In Jon Barnes, Alejna Brugos, Stefanie Shattuck-Hufnagel and Nanette Veilleux (eds.), *Proceedings of Speech Prosody, Special Session Satellite Session on Framing speech: Celebrating 40 Years of Inquiry with Stefanie Shattuck-Hufnagel*, Boston University, 1250–1253.
- West, Robert, Hristo Paskov, Jure Leskovec and Chris Potts (2014) Exploiting social network structure for person-to-person sentiment analysis. *Transactions of the ACL* 2(2), 297–310.
- White, Peter R.R. (2015) Appraisal Theory, in Karen Tracy, Cornelia Ilie and Todd Sandel (eds.), *The International Encyclopedia of Language and Social Interaction*, Hoboken: Wiley-Blackwell.
- Wiebe, Janyce, Theresa Wilson, Rebecca Bruce, Matthew Bell and Melanie Martin (2004) Learning subjective language. *Computational Linguistics* 30(3), 277–308.
- Wiltschko, Martina, Derek Denis and Alexandra d’Arcy (2018) Deconstructing variation in pragmatic function: A transdisciplinary case study. *Language in Society* 47(4), 569–599.
- Zillig, Werner (1982) *Bewerten: Sprechakttypen der bewertenden Rede*. Tübingen: Niemeyer.