

GASC: Genre-Aware Semantic Change for Ancient Greek

Valerio Perrone

Amazon, Berlin*

vperrone@amazon.com

Marco Palma

University of Warwick

m.palma@warwick.ac.uk

Simon Hengchen

University of Helsinki

simon.hengchen@helsinki.fi

Alessandro Vatri

University of Oxford
The Alan Turing Institute

avatri@turing.ac.uk

Jim Q. Smith

University of Warwick
The Alan Turing Institute

j.q.smith@warwick.ac.uk

Barbara McGillivray

University of Cambridge
The Alan Turing Institute

bmcgillivray@turing.ac.uk

Abstract

Word meaning changes over time, depending on linguistic and extra-linguistic factors. Associating a word's correct meaning in its historical context is a central challenge in diachronic research, and is relevant to a range of NLP tasks, including information retrieval and semantic search in historical texts. Bayesian models for semantic change have emerged as a powerful tool to address this challenge, providing explicit and interpretable representations of semantic change phenomena. However, while corpora typically come with rich metadata, existing models are limited by their inability to exploit contextual information (such as text genre) beyond the document timestamp. This is particularly critical in the case of ancient languages, where lack of data and long diachronic span make it harder to draw a clear distinction between polysemy (the fact that a word has several senses) and semantic change (the process of acquiring, losing, or changing senses), and current systems perform poorly on these languages. We develop GASC, a dynamic semantic change model that leverages categorical metadata about the texts' genre to boost inference and uncover the evolution of meanings in Ancient Greek corpora. In a new evaluation framework, our model achieves improved predictive performance compared to the state of the art.

1 Introduction

Change and its precondition, variation, are inherent in languages. Over time, new words enter the lexicon, others become obsolete, and existing words acquire new senses. These changes are grounded in cognitive, social, and contextual factors, and can be realized in different ways. For example, in Old English *thing* meant 'a public assembly'¹

*Work done prior to joining Amazon.

¹In the remainder of this paper, we use *emphasis* to refer to a word and 'single quotes' for any of its senses.

and currently it more generally means 'entity'. Semantic change research has a number of practical applications, beyond historical linguistics research, including new sense detection in computational lexicography and information retrieval for historical texts that allows to restrict a search to certain word senses (e.g. the old sense of the English adjective *nice* as 'silly'). To take an example from recent semantic change in English, the verb *tweet* used to be uniquely associated with birds' sounds and has recently acquired a new sense related to the social media platform Twitter. However, in this as in many other cases, the original sense co-exists with the new one, and specific contexts or genres will select one over the other. This is known as synchronic variation, and can be successfully modelled probabilistically, as advocated by several authors (see e.g. [Jenset and McGillivray \(2017\)](#)). The close relationship between innovation and variation is well-known in historical linguistics, and critical to ancient languages. Indeed, the unavailability of balanced corpora due to the limited amount of data at our disposal makes it crucial for models to explicitly account for confounding variables like genre, so as to enable them to use all existing data.

To address these challenges, we introduce GASC (**Genre-Aware Semantic Change**), a novel dynamic Bayesian mixture model for semantic change. In this model, the evolution of word senses over time is based not only on distributional information of lexical nature, but also on additional features, specifically genre. This allows GASC to decouple sense probabilities and genre prevalence, which is critical with genre-unbalanced data such as ancient languages corpora. The value of incorporating genre information in the model goes beyond literary corpora and historical language data and can be applied to recent data spanning over a period of time where text type information is critical, for example in specialized domains. Explicitly mod-

elling genres also makes it possible to address a number of additional questions, revealing the genre most likely associated to a given sense, the most unusual sense for a genre, and which genres have the most similar senses. Naturally, this framework can be applied to different categorical metadata about the text, such as author, geography, or style.

Ancient Greek is an insightful test case for several reasons. First, Ancient Greek words tend to have a particularly high number of senses (Clarke, 2010), and Ancient Greek texts display a large number of literary genres. Second, we can use data spanning several centuries. Third, Ancient Greek scholarship provides high-quality data to validate automatic systems. Top-quality transcribed Ancient Greek texts are available, eliminating the need for OCR correction. Finally, polysemous words are particularly sensitive to register variation and the distribution of senses can vary greatly across registers (Leiwo et al., 2012). As most extant texts are literary and relatively conservative from a linguistic perspective, we expect genre (the type of a text) and register (the fact that different varieties of language are used in particular situations) to play a significant role in the variation of sense distributions in polysemous words. The word *mus*, for instance, can mean ‘mouse’, ‘muscle’, or ‘mussel’. The effect of genre on the distribution of its meaning can be estimated visually from Figure 1. In this graph, lines represent the percentage of the occurrences of the target word in a literary genre across centuries, while bars represent the percentage of the occurrences of a specific sense of the target word across centuries. If any line shows a similar trend to that of any set of bars, we may estimate that genre might play a more decisive role than diachrony in determining variation in the distribution of senses. Here, the distribution of ‘muscle’ over time (pink bars) closely follows the distribution of this word in technical genres over time (blue line), suggesting that the effect of genre should be incorporated into semantic change models.

2 Related work

Semantic change in historical languages, especially on a large scale and over a long time period, is an under-explored, but impactful research area. Previous work has mainly been qualitative in nature, due to the complexity of the phenomenon (cf. e.g. Leiwo et al. (2012)). In recent years, NLP research has made great advances in the area of

semantic change detection and modelling (for an overview of the NLP literature, see Tang (2018) and Tahmasebi et al. (2018)), with methods ranging from topic-based models (Boyd-Graber et al., 2007; Cook et al., 2014; Lau et al., 2014; Wijaya and Yeniterzi, 2011; Frermann and Lapata, 2016), to graph-based models (Mitra et al., 2014, 2015; Tahmasebi and Risse, 2017), and word embeddings (Kim et al., 2014; Basile and McGillivray, 2018; Kulkarni et al., 2015; Hamilton et al., 2016; Dubossarsky et al., 2017; Tahmasebi, 2018; Rudolph and Blei, 2018; Jatowt et al., 2018; Dubossarsky et al., 2019). However, such models are purely based on words’ lexical distribution information and do not account for language variation features such as text type because genre-balanced corpora are typically used.

With the exception of Bamman and Crane (2011) and Rodda et al. (2017), no previous work has focussed on ancient languages. Recent work on languages other than English is rare but exists: Falk et al. (2014) use topic models to detect changes in French and Hengchen (2017) uses similar methods to tackle Dutch. Cavallin (2012) and Tahmasebi (2018) focus on Swedish, with the comparison of verb-object pairs and word embeddings, respectively. Zampieri et al. (2016) use SVMs to assign a time period to text snippets in Portuguese, and Tang et al. (2016) work on Chinese newspapers using S-shaped models. Most work in this area focusses on simply detecting the occurrence of semantic change, while Frermann and Lapata (2016)’s system, SCAN, takes into account synchronic polysemy and models how the different word senses evolve across time.

Our work bears important connections with the topic model literature. The idea of enriching topic models with document-specific author meta-data was explored in Rosen-Zvi et al. (2004) for the static case. Several time-dependent extensions of Bayesian topic models have been developed, with a number of parametric and nonparametric approaches (Blei and Lafferty, 2006; Rao and Teh, 2009; Ahmed and Xing, 2012; Dubey et al., 2013; Perrone et al., 2017). In this paper, we transfer such ideas to semantic change, where each datapoint is a bag of words associated to a single sense (rather than a mixture of topics). Excluding cases of intentional ambiguity, which we expect to be rare, we can safely assume that there are generally no ambiguities in a context, and each word instance

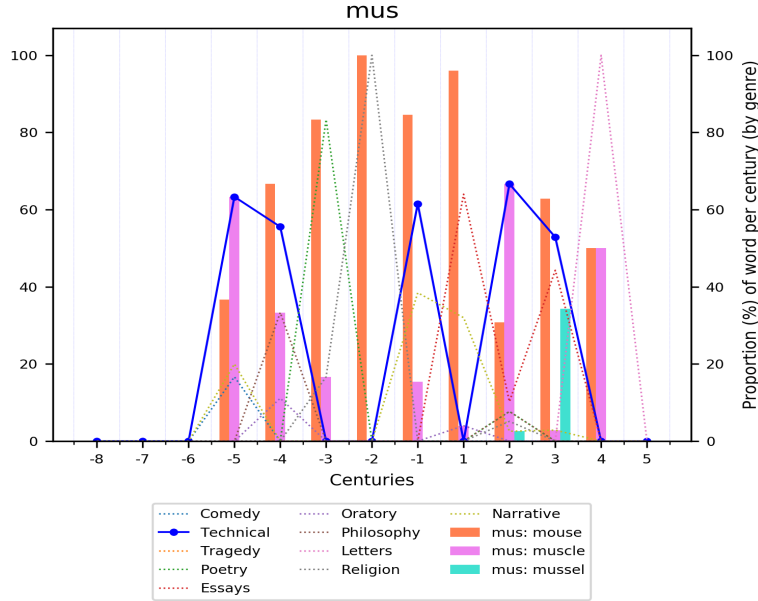


Figure 1: Distribution of *mus* ‘mouse’/‘muscle’/‘mussel’ by genre vs its senses over time. Lines track *mus* proportions in each genre and century, while bars show the *mus* occurrence proportions with each sense and century.

maps to a single sense.

3 The model

We start with a lemmatized corpus pre-processed into a set of text snippets, each containing an instance of the word under study (referred to as “target word” in the remainder). Each snippet is a fixed-sized window W of 5 words to the left and right of the target word. The inferential task is to detect the sense associated to the target word in the given context, and describe the evolution of sense proportions over time.

The generative model for GASC is presented in Algorithm 1 and illustrated by the plate diagram in Figure 2. First, suppose that throughout the corpus the target word is used with K different senses, where we define a sense at time t as a distribution ψ_k^t over words from the dictionary. These distributions are used to generate text snippets by drawing each of their words from the dictionary based on a Multinomial distribution (line 13 in Algorithm 1). Based on the intuition that each genre is more or less likely to feature a given sense, we assume that each of G possible text genres determines a different distribution over senses (lines 3-4). Each observed document snippet is then associated with a genre-specific distribution over senses $\phi_{g^d}^t$ at time t , where g^d is the observed genre for document d . Crucially, conditioning on the observed genre we have a specific distribution over senses account-

ing for genre-specific word usage patterns (line 11). On the other hand, to make sure senses can be uniquely identified across genres, we associate each sense to the same probability distribution over words for all genres. We let word (line 7) and sense distributions (line 4) evolve over time with Gaussian changes, ensuring smooth transitions. The coupling between sense probabilities over time is controlled by K^ϕ , the sense probability precision parameter, so that the larger K^ϕ , the stronger the coupling between the sense probabilities over time. We place a Gamma prior over K^ϕ with hyperparameters a and b (line 1), and infer K^ϕ from the data. We fix K^ψ , the word probability precision parameter.

Hyperparameter settings The model can be applied to different inferential goals: we can focus on the evolution of sense probabilities or on the changes within each sense. For each of these aims, we can use several hyperparameter combinations for K^ϕ , which is drawn from the prior distribution as determined by a and b , and K^ψ . Specifically, we consider the following 3 settings. Setting 1: $a = 7, b = 3, K^\psi = 10$, as in [Fremann and Lapata \(2016\)](#). Setting 2: $a = 7, b = 3, K^\psi = 100$. This aims at enforcing less variation within senses over time. Setting 3: $a = b = 1, K^\psi = 100$. This still keeps the bag of words stable for each sense, but also induces less smoothing for sense probabilities over time. Setting 3 allows probabilities

to vary widely across centuries. We also expect a large K^ψ to reduce the likelihood of dramatic changes within the same sense across contiguous time periods, and to favour the emergence of new senses. If not otherwise specified, we use setting 3. Other settings (like setting 3 with $K^\psi = 10$) are not recommended since allowing relevant changes over time both in sense probabilities and bag of words might harm interpretability. A final parameter is the window size W , namely the number of words surrounding an instance of the target. While larger windows increase the range of captured dependencies, noise can be introduced in the form of irrelevant contextual words. As in SCAN, we fixed the window size W to 5 for all methods.

Algorithm 1: GASC generative model

```

1 Draw  $K^\phi \sim \text{Gamma}(a, b)$ ;
2 for time  $t = 1, \dots, T$  do
3   for genre  $g = 1, \dots, G$  do
4     Draw sense distribution  $\phi_g^t \mid$ 
        $\phi_g^{-t}, K^\phi \sim N(\frac{1}{2}(\phi_g^{t-1} + \phi_g^{t+1}), K^\phi)$ 
5   end
6   for sense  $k = 1, \dots, K$  do
7     Draw word distribution
        $\psi_k^t \mid \psi^{-t}, K^\psi \sim$ 
        $N(\frac{1}{2}(\psi_k^{t-1} + \psi_k^{t+1}), K^\psi)$ 
8   end
9   for document  $d = 1, \dots, D_t$  do
10    Let  $g^d$  be the observed genre;
11    Draw sense
        $z^d \mid g^d \sim \text{Mult}(\text{softmax}(\phi_{g^d}^t))$ ;
12    for context position  $i = 1, \dots, W$  do
13      Draw word
        $w^{d,i} \sim \text{Mult}(\text{softmax}(\psi^{t,z^d}))$ ;
14    end
15  end
16 end

```

Posterior inference For posterior inference, we extend the blocked Gibbs sampler proposed in [Frumm and Lapata \(2016\)](#). The full conditional is available for the snippet-sense assignment, while to sample the sense and word distributions we adopt the auxiliary variable approach from [Mimno et al. \(2008\)](#). The sense precision parameters are drawn from their conjugate Gamma priors. For the distribution over genres we proceed as follows. First, sample the distribution over senses ϕ_g^t for each genre $g = 1, \dots, G$ following [Mimno et al. \(2008\)](#).

Then, sample the sense assignment conditioned on the observed genre from its full conditional: $p(z^d \mid g^d, \mathbf{w}, t, \phi, \psi) \propto p(z^d \mid g^d, t)p(\mathbf{w} \mid t, z^d) = \phi_g^t \prod_{w \in \mathbf{w}} \psi_w^{t,z^d}$. This setting easily extends to sample genre assignments for tasks where, for example, some genre metadata are missing.

4 Evaluation framework

Evaluating models tackling lexical semantic change is notoriously challenging. Frameworks are either lacking or focus on very specific types of sense change ([Schlechtweg et al., 2018](#); [Tahmasebi et al., 2018](#)). Exceptions are [Kulkarni et al. \(2015\)](#), [Basile and McGillivray \(2018\)](#) and [Hamilton et al. \(2016\)](#), who focus on the change points of word senses. However, in the case of Ancient Greek (and other historical languages), corpora typically contain gaps and uneven distribution of text genres, and semantic change is so closely related to polysemy that it is hard to find a specific point in time when a new sense emerged in the language. Therefore, it is more appropriate to take a probabilistic approach to model sense distribution, and devise an evaluation approach that fits this. Although historical dictionaries and traditional philology describe the evolution of word senses over time, they do not necessarily reflect the evidence from corpora on which models can be evaluated, and often only provide insights into the appearance of a new sense, rather than the relative predominance of a word’s senses across time. These reasons led us to craft a novel evaluation dataset and framework, which reflects the data on which the model is evaluated, and allows for a finer-grained evaluation of the predominance of word senses across time.

4.1 Ancient Greek corpus

We used the Diorisis Annotated Ancient Greek Corpus ([Vatri and McGillivray, 2018](#)), consisting of 10,206,421 lemmatized and part-of-speech-tagged words. The corpus contains 820 texts spanning between the beginnings of the Ancient Greek literary tradition (8th century BC) and the 5th century AD. The corpus covers a number of Ancient Greek literary and technical genres: poetry (narrative, choral, epigrams, didactic), drama (tragedy, comedy), oratory, philosophy, essays, narrative (historiography, biography, mythography, novels), geography, religious texts (hymns, Jewish and Christian Scriptures, theology, homilies), technical literature (medicine, mathematics, natural science,

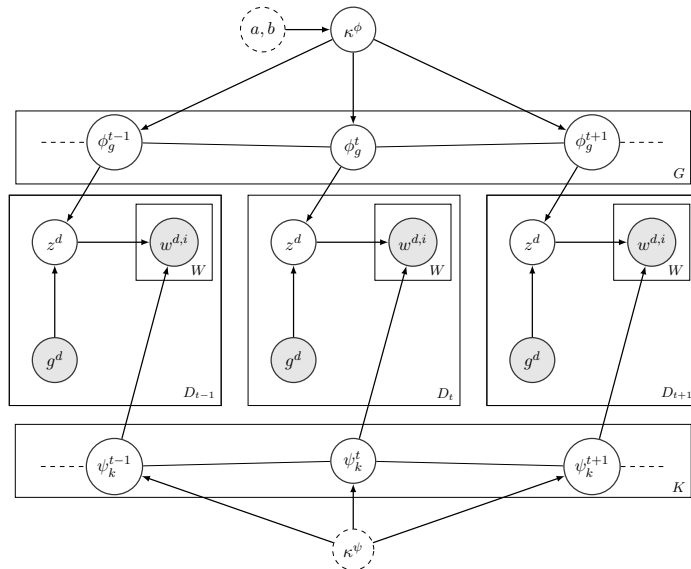


Figure 2: GASC plate diagram with 3 time periods.

tactics, astronomy, horsemanship, hunting, politics, art history, rhetoric, literary criticism, grammar, and letters. In technical texts, we expect polysemous words to have a technical sense. On the other hand, in works more closely representing general language (comedy, oratory, historiography) we expect words to appear in their more concrete and less metaphorical senses; we cannot assume that this distribution holds in a number of other genres such as philosophy and tragedy. Whilst genre-annotated corpora are not especially common in NLP, where most tasks rely on specific genres (e.g. Twitter) or on genre-balanced corpora such as COHA (Davies, 2002), they are more prevailing within humanities, and especially classics. Additionally, research on automated genre identification has been flourishing for decades (e.g. Kessler et al. (1997)), making the need for genre information in a potential corpus not as much of a hindrance as can be thought.

4.2 Log-likelihood evaluation

First, we compared GASC with the state-of-the-art (SCAN) in terms of held-out data log-likelihood. We chose 50 targets that could be identified as polysemous (e.g. the verb *legō*, whose senses are ‘gather’ and ‘tell’) based on two criteria: high frequency and a suitably clear-cut range of meanings. We initially based our selection on the secondary literature and chose 17 words from the well-studied vocabulary of Ancient Greek aesthetics (Pollitt, 1974). We complemented this selection with the inclusion of the 33 most frequent clearly polysemous words identified by an Ancient Greek expert

in a frequency-ranked word list extracted from the Diorisis corpus. The necessity to identify manually suitable words led us to limit their number to 50. For each one of these target words, we randomly divided the corpus into a train (80%) and test set (20%). Results on the 50-word dataset are in Section 5.

4.3 Expert annotation

To evaluate our method against ground truth, we proceeded as follows. We selected three three target words (*mus* ‘mouse’/‘muscle’/‘mussel’ and *harmonia* ‘fastening’/‘agreement’/‘musical scale, melody’, and *kosmos* ‘order’/‘world’/‘decoration’) based on their frequency and clear-cut polysemy, as indicated by the standard scholarly Ancient Greek-English dictionary (Liddell et al., 1996) and traditional philological scholarship on their semantics (Pollitt 1974 on *harmonia* and *kosmos*). These words are especially suitable for an exploratory case study because they exhibit an abstract sense and a concrete counterpart in general, non-technical vocabulary, and are attested in most of the time periods covered by the corpus and across different literary genres. Two Ancient Greek experts manually annotated the whole corpus by tagging the senses of the target words in context. One expert selected the correct sense for each occurrence of *mus* and *harmonia*, and the other expert performed the same task on *kosmos*. The results of each expert’s annotation task were not reviewed by the other expert (Vatri et al. 2019 for the dataset). Table 1 shows an example from the

annotated dataset for the word *kosmos*. The annotators also marked when the semantic annotation was purely based on the target word context, which is the evidence on which the model can rely (category ‘collocates’). Only annotations based on collocates were retained in the evaluation. Using this information, the relative frequency of each sense for each target word in any time slice becomes computable, and was used to create ground-truth data on the diachronic predominance of a word’s senses as reflected in the corpus.

4.4 Automatic sense labelling

For every time period T , inferred sense k , and genre G , GASC outputs a distribution of words with associated probabilities. For instance, the output for *kosmos* (‘order’, ‘world’, or ‘decoration’) in oratory at time 0 includes:

```
T=0,k=0:
aêr (0.069); mousikos (0.059); gê (0.056); harmonia
(0.034); ouranos (0.033); logos (0.030);
gignomai (0.021); sphaira (0.021); pselion
(0.020); apaiteô (0.019);
T=0,k=1:
polis (0.035); asebeia (0.014); politeia (0.012);
proteros (0.012); naus (0.012); pentêkonta
(0.011); aei (0.011); hama (0.011); peripeteia
(0.011); oikia (0.011).
```

These distributions can be interpreted by experts based on the meanings of the words they group and thus associated to the senses of the target word. Here, $K = 0$ includes *aêr* (‘air’), *gê* (‘earth’), *ouranos* (‘sky’), and *sphaira* (‘sphere, globe’), which point to the meaning of *kosmos* as ‘world’. The list for $K = 1$ includes *polis* (‘city’), *asebeia* (‘impiety’), *politeia* (‘constitution’), and *oikia* (‘household’), which point to the meaning of *kosmos* as ‘order’. On the other hand, the expert annotation provides lists of corpus occurrences of the target word, each associated to a sense label. In Table 1, the sense label is ‘kosmos-world’ and we can associate lemmas such as *ouranos* ‘sky’ and *sphairoeides* ‘spherical’ to this sense, as these lemmas occur in the corpus context of this target word.

To evaluate against expert annotation, we automatically match the word senses assigned by the annotators (denoted by s) with the senses outputted by the model (denoted by k). To achieve this, we first measured how closely each model sense k matches each expert sense s . We assigned a confidence score to every possible (k, s) pair by comparing the words associated to k in the model output and the words co-occurring with the target word in the annotated corpus sentences labelled with the expert-assigned sense s . For *kosmos* with $k = 0$,

we compare words from the model output, such as *ouranos* ‘sky’, *gê* ‘earth’, and *sphaira* ‘sphere’ with words from the context of the annotated sentences, such as *sphairoeides* ‘spherical’ and *ouranos* ‘sky’. We then considered two elements. For words from the model output, we consider the normalized probability with which these words w_i are associated to the model sense k , i.e. $P(w_i|k)$. For *kosmos*, *aêr* ‘air’ is associated to probability 0.069, *gê* ‘earth’ to 0.056, and *ouranos* ‘sky’ to 0.033. For the context words from the annotated data, we consider the degree to which these words are associated to an expert sense. In the example of *kosmos* from Table 1, this is calculated based on how many different senses a context word like *ouranos* ‘sky’ or *sphairoeides* ‘spherical’ is associated to. To measure this degree of association we define the expert score $m(w_i, s)$ of word w_i as 1 divided by the number of senses assigned by the experts to this word. If the word is associated to only one sense s in the annotated data, its expert score $m(w_i, s)$ will be highest (1); if it is associated to two senses, its expert score is 0.5; if it is not assigned to the sense s by the experts, its expert score $m(w_i, s)$ is 0. Formally, we define the confidence score of a pair of model sense k and expert-assigned sense s as $\text{conf}(k, s) = \sum_i P(w_i|k) * m(w_i, s)$. The score is highest when $P(w_i)$ and $m(w_i, s)$ are highest for all words. In extreme cases, $P(w_i)$ will be 1 if the model estimated w_i to be associated to sense k with probability 1 and $m(w_i, s)$ is 1 (i.e. w_i is only found in contexts labelled as s by the experts). This points to k and s being associated to the same words, and thus being the same sense. The confidence is lowest when k and s do not share words, in which case either $P(w_i) = 0$ or $m(w_i, s) = 0$. In contrast with clustering overlap techniques like purity or rand index, we ensure words with a higher inferred probability and uniquely associated to a sense weigh more. The confidence scores were used to find the best matching pair (k, s) : for every expert sense s we selected the sense(s) k for which $\text{conf}(k, s)$ was higher than the random baseline (1 over the number of expert senses) and higher than the sum of the 2d and 3rd best confidence scores, when possible, or assigned NA when both conditions were not matched. We consider NA as an additional expert sense whenever the expert assigned a sense based on other factors than lexical context.

After matching inferred and expert-assigned

date	genre	author	work	target word	sense id
-335	Technical	Aristotle	De Mundo	kosmos	kosmos:world

Table 1: Example from annotated dataset displaying *Tou de sumpantos ouranou te kai kosmou sphairoeidous ontos kai kinoumenou kathaper eipon* (“The whole of the heaven, the whole cosmos, is spherical, and moves continuously, as I have said”), containing the target *kosmos* and its expert-assigned sense ‘world’ (date: 335 B. C.).

senses, we computed precision and recall. For every target word and matched pair (s, k) , a word is considered correctly assigned to sense k if it also appeared within a 5-word window of the target word in the expert annotation for s . In the example above for *kosmos*, $k = 0$ and $s = \text{‘kosmos-world’}$, one such word is *ouranos* ‘sky’ as it appears in the model output for $k = 0$ and in the context window of a sentence labelled as ‘kosmos-world’ by the annotators. Moreover, we weighted each word by the inferred probability to account for the different degrees of association of words to senses. Specifically, we defined precision as the ratio between the number of words correctly assigned to k , weighted by their respective normalised model-estimated probabilities, and the number of words assigned to k by the model. This metric is based on the distributional hypothesis whereby words occurring in similar contexts tend to exhibit similar meanings. We computed precision after stop word removal, limiting the noise from uninformative contextual words. We defined precision in terms of the words assigned to a sense also appearing within a 5-word window of the target in the expert annotation. Our model, as SCAN, only considers those context words to determine word senses, and for the ground truth evaluation we only retained the cases in which the annotators could disambiguate words based purely on their context. We defined recall as the ratio between the number of all words correctly assigned to k (weighted by their probabilities) and number of words assigned to sense s by the experts (weighted by their expert scores). For each model, precision and recall scores for each (s, k) pair were averaged and used as final scores. Since recall directly depends on the number of expert words, the metric can only be used to compare models for a specific target word. While the proposed assessment focusses on dynamic mixture models, it can be generalised to any probabilistic model by considering the posterior probability of the gold word sense.

5 Experiments

Predictions on held-out words Considering the 50-word dataset described in Section 4, we evaluated the predictive performance in terms of log-likelihood of held-out data for SCAN (not using any genre information), GASC-all (GASC with all the $G = 10$ available genres) and GASC-narr (GASC with 2 genres, Narrative vs. non Narrative). Narrative and Technical are such that all 50 words occurred at least once in the training and test sets, and analogous results are obtained when GASC with Technical vs. non Technical. For each model, we compared the 3 hyperparameter settings previously reported, with higher scores indicating that a model is better at explaining unseen data.

Figure 3 shows predictive log-likelihood scores for a range of K , with the results averaged over 50 leave-one-out folds. Each time, the scores were averaged under the final 10 samples of the latent variables, out of 1000 MCMC iterations. On average, GASC-narr consistently outperforms SCAN across every K and hyperparameter setting. On the other hand, SCAN has a higher held-out log-likelihood than GASC-all. Exploiting some information on the genre yields better predictions, while using all genres attested in the corpus is not effective as some genres are not sufficiently represented by the data. Figure 3 also shows that the best predictions over unseen data are obtained for K between 10 and 15. Higher K values tend to introduce noisy senses with no improvement for the model output. In addition, Setting 3 worked better or on par with other settings. In the next section, we fix the hyperparameters and use a validation set of words that were not part of the 50 targets of this experiment.

Ground truth recovery We explored the ability to recover ground truth when available. For *mus*, experts annotated 205 instances, of which 198 were assigned to one of the 3 senses ‘mouse’, ‘mussel’, and ‘muscle’; out of these 198 assignments, 114 were based on lexical contextual information only (category ‘collocates’) and were retained for the evaluation. For *harmonia*, the number of annotated occurrences was 599, of which 411 were of the type

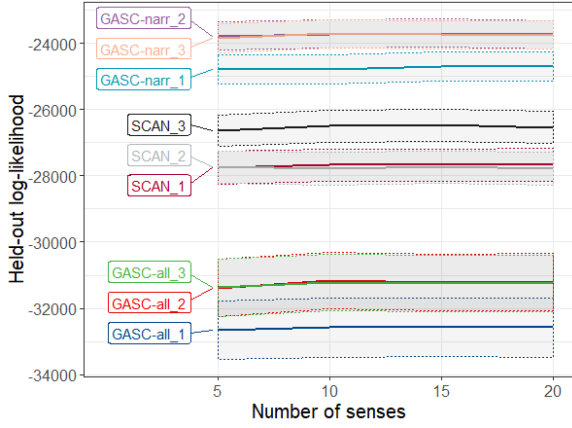


Figure 3: Held-out mean log-likelihood varying K (the larger, the better). Shaded areas are ± 1 standard error.

harmonia	'agreement, harmony'	Technical ($\rho = 0.888$, $p < 0.0001$) Narrative ($\rho = 0.719$, $p = 0.006$) Essays ($\rho = 0.561$, $p = 0.046$)
	'fastening'	Narrative ($\rho = 0.663$, $p = 0.013$)
	'stringing, music scale'	Technical ($\rho = 0.817$, $p = 0.001$) Philosophy ($\rho = 0.632$, $p = 0.02$) Essays ($\rho = 0.598$, $p = 0.031$)
kosmos	'decoration'	Narrative ($\rho = 0.887$, $p = 0.001$) Technical ($\rho = 0.705$, $p = 0.023$) Oratory ($\rho = 0.664$, $p = 0.036$)
	'order'	Technical ($\rho = 0.875$, $p = 0.001$) Narrative ($\rho = 0.862$, $p = 0.001$)
	'world'	Technical ($\rho = 0.792$, $p = 0.006$) Oratory ($\rho = 0.723$, $p = 0.018$)
mus	'mouse'	Narrative ($\rho = 0.813$, $p = 0.001$) Essays ($\rho = 0.743$, $p = 0.004$)
	'muscle'	Technical ($\rho = 0.766$, $p = 0.002$)
	'mussel'	Narrative ($\rho = 0.736$, $p = 0.004$) Essays ($\rho = 0.736$, $p = 0.004$) Poetry ($\rho = 0.613$, $p = 0.026$)

Table 2: Correlations between senses and genres for manually annotated target words.

'collocates'. For *kosmos*, 1,411 occurrences were annotated, of which 1,406 were assigned to a sense, and in 1,102 cases the annotation was of the type 'collocates'. Only the annotations of the type 'collocates' were kept for the expert sense distribution, and thus for the evaluation. We identified genres with the largest effect on the distribution of senses by the Spearman's Rank Correlation Coefficient for each word-sense s between the frequency $f(s)$ of s across centuries and the frequency $f(s, g)$ of s in each genre g across centuries (Table 2). Significant correlation between $f(s)$ and any $f(s, g)$ suggests that variation in the frequency of a word sense across centuries is not due to diachronic change, but to how frequently s is attested in g in each century (and, ultimately, to the amount of texts representing g in each century). Given the amount of available data and the size of the correlations, we considered the genres Technical and non-Technical for *mus* and *harmonia*, and both Technical and

non-Technical and Narrative and non-Narrative for *kosmos*. These words were selected as examples of polysemous words (a) with a range of clearly distinct senses (such as 'mus', whose three senses are strikingly diverse), (b) attested in most, if not all, the time periods covered by the corpus, and (c) attested across a number of genres. As expert annotations of semantic change in Ancient Greek corpora are virtually unavailable, this choice also allowed us to leverage ground truth for validation.

We compared SCAN with GASC and GASC-independent, a simpler version that fits independent models to sets of documents sharing the same genre, so that parameters and senses are inferred independently across genres (while in GASC senses are shared but their probability distributions are independent across genres). First, we compared word senses across time with expert-annotated data. Figure 4 shows the time distribution of the senses of *kosmos* in the expert annotation (top) and as outputted by SCAN and GASC run on Narrative vs. non-Narrative (bottom). For non-narrative texts, the GASC sense distribution successfully captures the 'world' sense arising only after 400 BC, which is less clear for SCAN. Second, we computed precision, recall, and F1 scores (the harmonic mean of precision and recall) to determine how closely the words assigned to a sense match the ones assigned by experts (Table 3). For GASC, the values average precision, recall, and F1-score for {Technical, non-Technical} for *mus* and *harmonia* and {Narrative, Non-Narrative} for *kosmos*. The results show that, for the most represented targets, genre information improves the ability to recover the ground truth.

6 Conclusion

We introduced GASC, a Bayesian model to study the evolution of word senses in ancient texts. We performed this analysis conditional on the text genre, demonstrating that the ability to harness genre metadata addresses a fundamental challenge in disambiguating word senses in ancient Greek. In experiments we showed that GASC provides interpretable representations of the evolution of word senses, and achieves improved predictive performance compared to the state of the art. Further, we established a new framework to assess model accuracy against expert judgement. To our knowledge, no previous work has systematically compared the estimates from a statistical model to manual semantic annotations of ancient texts.

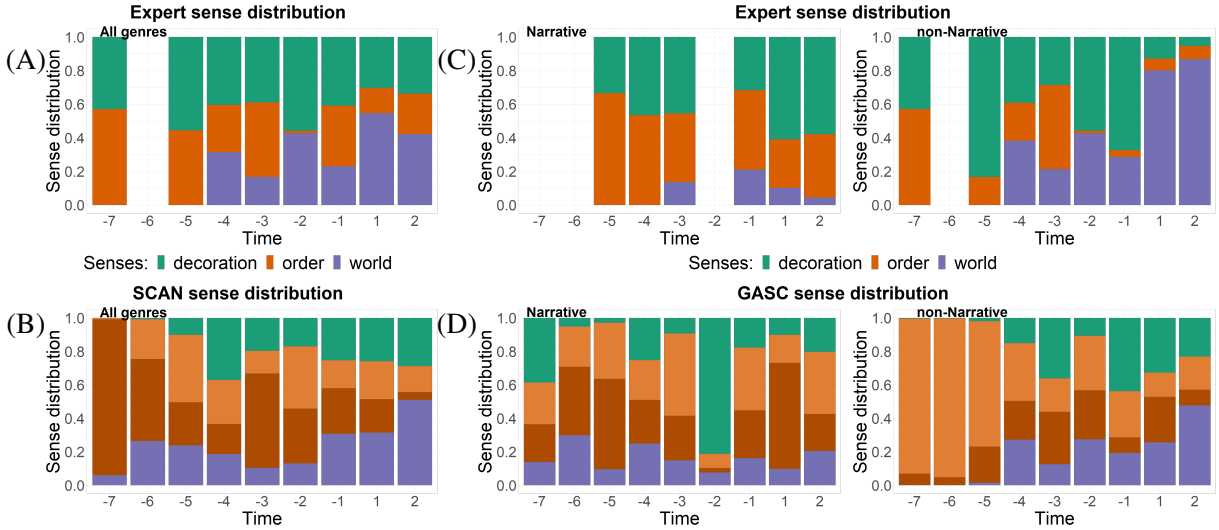


Figure 4: Expert annotation (top) vs SCAN and GASC (bottom). Each stacked bar represents all *kosmos* occurrences in a given time. Colours denoting senses are matched between plots. Both shades of orange map to ‘order’, but the fourth sense in (B) and (D) is NA (i.e., $\text{conf}(k, s)$ not higher than the random baseline and not higher than the sum of 2nd and 3rd best confidence scores).

Word/Model	SCAN			GASC-independent			GASC		
	P	R	F1	P	R	F1	P	R	F1
<i>mus</i>	0.430	0.477	0.452	0.420	0.442	0.431	0.224	0.298	0.253
<i>harmonia</i>	0.527	0.708	0.603	0.582	0.729	0.646	0.497	0.481	0.484
<i>kosmos</i>	0.405	0.586	0.478	0.362	0.447	0.399	0.525	0.611	0.595

Table 3: SCAN vs GASC on *mus* (‘mouse’, ‘muscle’, ‘mussel’), *harmonia* (‘abstract’, ‘concrete’, ‘musical’), and *kosmos* (‘order’, ‘decoration’, ‘world’) in terms of precision (‘P’), recall (‘R’), and F1-score (‘F1’).

This work can be seen as a step towards the development of richer evaluation schemes and models that can embed expert judgement. Future work could encode more structured cross-genre dependencies, or allow for change points that occur in the light of exogenous forces by historical events.

Acknowledgements

This work was supported by The Alan Turing Institute under the EPSRC grant EP/N510129/1 and the seed funding grant SF042. We would also like to thank Viivi Lähteenoja for her work on the annotation of the word *kosmos*, Dr Lea Frermann for her advice, Prof Geoff Nicholls for his insights on the model, as well as the anonymous reviewers for their helpful comments.

References

Amr Ahmed and Eric P. Xing. 2012. Timeline: A dynamic hierarchical Dirichlet process model for recovering birth/death and evolution of topics in text stream. *UAI*, abs/1203.3463.

David Bamman and Gregory Crane. 2011. Measuring

Historical Word Sense Variation. *Proceedings of the 11th ACM/IEEE-CS Joint Conference on Digital Libraries*, pages 1–10.

Pierpaolo Basile and Barbara McGillivray. 2018. *Discovery Science*, volume 11198 of *Lecture Notes in Computer Science*, chapter Exploiting the Web for Semantic Change Detection. Springer-Verlag.

David M. Blei and John D. Lafferty. 2006. Dynamic topic models. *ICML*, pages 113–120.

Jordan Boyd-Graber, David M. Blei, and Xiaojin Zhu. 2007. A topic model for word sense disambiguation. *EMNLP*.

Karin Cavallin. 2012. Automatic extraction of potential examples of semantic change using lexical sets. In *KONVENS*, pages 370–377.

Michael Clarke. 2010. Register Variation. In Egbert J. Bakker, editor, *A companion to the Ancient Greek language*, pages 120–33. Wiley-Blackwell, Chichester/Malden, Mass.

Paul Cook, Jey Han Lau, Diana McCarthy, and Timothy Baldwin. 2014. Novel word-sense identification. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 1624–1635.

- Mark Davies. 2002. *The Corpus of Historical American English (COHA): 400 million words, 1810-2009*. Brigham Young University.
- Avinava Dubey, Ahmed Hefny, Sinead Williamson, and Eric P. Xing. 2013. A Nonparametric Mixture Model for Topic Modeling over Time. *SDM*, pages 530–538.
- Haim Dubossarsky, Simon Hengchen, Nina Tahmasebi, and Dominik Schlechtweg. 2019. Time-out: Temporal referencing for robust modeling of lexical semantic change. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Florence, Italy. Association for Computational Linguistics.
- Haim Dubossarsky, Daphna Weinshall, and Eitan Grossman. 2017. Outta control: Laws of semantic change and inherent biases in word representation models. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1136–1145.
- Ingrid Falk, Delphine Bernhard, and Christophe Gérard. 2014. De la quenelle culinaire à la quenelle politique : identification de changements sémantiques à l’aide des Topic Models. In *21ème conférence sur le Traitement Automatique des Langues Naturelles*, page 443.
- Lea Frermann and Mirella Lapata. 2016. A Bayesian model of diachronic meaning change. *Transactions of the Association for Computational Linguistics*, 4:31–45.
- William L Hamilton, Jure Leskovec, and Dan Jurafsky. 2016. Diachronic word embeddings reveal statistical laws of semantic change. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, volume 1, pages 1489–1501.
- Simon Hengchen. 2017. *When Does it Mean?: Detecting Semantic Change in Historical Texts*. Ph.D. thesis, Université libre de Bruxelles.
- Adam Jatowt, Ricardo Campos, Sourav S Bhowmick, Nina Tahmasebi, and Antoine Doucet. 2018. Every word has its history: Interactive exploration and visualization of word sense evolution. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pages 1899–1902. ACM.
- Gard B. Jensen and Barbara McGillivray. 2017. *Quantitative Historical Linguistics. A Corpus Framework*. Oxford University Press, Oxford.
- Brett Kessler, Geoffrey Numberg, and Hinrich Schütze. 1997. Automatic detection of text genre. In *Proceedings of the eighth conference on European chapter of the Association for Computational Linguistics*, pages 32–38. Association for Computational Linguistics.
- Yoon Kim, Yi-I Chiu, Kentaro Hanaki, Darshan Hegde, and Slav Petrov. 2014. Temporal analysis of language through neural language models. *ACL 2014*, page 61.
- Vivek Kulkarni, Rami Al-Rfou, Bryan Perozzi, and Steven Skiena. 2015. Statistically significant detection of linguistic change. In *Proceedings of the 24th International Conference on World Wide Web*, pages 625–635. International World Wide Web Conferences Steering Committee.
- Jey Han Lau, Paul Cook, Diana McCarthy, Spandana Gella, and Timothy Baldwin. 2014. Learning word sense distributions, detecting unattested senses and identifying novel senses using topic models. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 259–270.
- Martti Leiwo, Hilla Halla-aho, and Marja Vierros. 2012. *Variation and Change in Greek and Latin. Papers and monographs of the Finnish Institute at Athens*. Foundation of the Finnish Institute at Athens, Helsinki.
- Henry George Liddell, Robert Scott, Henry Stuart Jones, and Roderick McKenzie. 1996. *A Greek-English lexicon*, 9th ed. edition. Clarendon Press Oxford University Press, Oxford New York.
- David Mimno, Hanna Wallach, and Andrew McCallum. 2008. Gibbs sampling for logistic normal topic models with graph-based priors. In *NIPS workshop*.
- Sunny Mitra, Ritwik Mitra, Suman Kalyan Maity, Martin Riedl, Chris Biemann, Pawan Goyal, and Animesh Mukherjee. 2015. An automatic approach to identify word sense changes in text media across timescales. *Natural Language Engineering*, 21(5):773–798.
- Sunny Mitra, Ritwik Mitra, Martin Riedl, Chris Biemann, Animesh Mukherjee, and Pawan Goyal. 2014. That’s sick dude!: Automatic identification of word sense change across different timescales. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1020–1029.
- Valerio Perrone, Paul A. Jenkins, Dario Spanò, and Yee Whye Teh. 2017. Poisson random fields for dynamic feature models. *Journal of Machine Learning Research*, 18(127):1–45.
- Jerome Jordan Pollitt. 1974. *The ancient view of Greek art: criticism, history, and terminology*. Number 26 in Yale publications in the history of art. Yale University Press, New Haven. OCLC: 299700455.
- Vinayak Rao and Yee Whye Teh. 2009. Spatial normalized gamma processes. In *Advances in neural information processing systems*, pages 1554–1562.

- Martina A. Rodda, Marco S.G. Senaldi, and Alessandro Lenci. 2017. Panta Rei: Tracking Semantic Change with Distributional Semantics in Ancient Greek. *Italian Journal of Computational Linguistics*, 3:11–24.
- Michal Rosen-Zvi, Thomas Griffiths, Mark Steyvers, and Padhraic Smyth. 2004. The author-topic model for authors and documents. *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*, pages 487–494.
- Maja Rudolph and David Blei. 2018. Dynamic embeddings for language evolution. In *Proceedings of the 2018 World Wide Web Conference on World Wide Web*, pages 1003–1011.
- Dominik Schlechtweg, Sabine Schulte im Walde, and Stefanie Eckmann. 2018. Diachronic usage relatedness (DUREL): A framework for the annotation of lexical semantic change. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics, NAACL 2018*.
- Nina Tahmasebi. 2018. A study on word2vec on a historical Swedish newspaper corpus. In *CEUR Workshop Proceedings. Vol. 2084. Proceedings of the Digital Humanities in the Nordic Countries 3rd Conference, Helsinki Finland, March 7-9, 2018*. University of Helsinki, Faculty of Arts.
- Nina Tahmasebi, Lars Borin, and Adam Jatowt. 2018. Survey of computational approaches to lexical semantic change detection. *Research Data Journal for the Humanities and Social Sciences*.
- Nina Tahmasebi and Thomas Risse. 2017. Finding individual word sense changes and their delay in appearance. In *Proceedings of the International Conference Recent Advances in Natural Language Processing, RANLP 2017*, pages 741–749.
- Xuri Tang. 2018. A state-of-the-art of semantic change computation. *arXiv preprint arXiv:1801.09872*.
- Xuri Tang, Weiguang Qu, and Xiaohe Chen. 2016. Semantic change computation: A successive approach. *World Wide Web*, 19(3):375–415.
- Alessandro Vatri, Viivi Lähteenoja, and Barbara McGillivray. 2019. Ancient Greek Semantic Change - Annotated Datasets and Code [dataset]. doi:10.6084/m9.figshare.c.4445420.
- Alessandro Vatri and Barbara McGillivray. 2018. The Diorisis Ancient Greek Corpus. *Research Data Journal for the Humanities and Social Sciences*.
- Derry Tanti Wijaya and Reyyan Yeniterzi. 2011. Understanding semantic change of words over centuries. In *Proceedings of the 2011 international workshop on DETecting and Exploiting Cultural diversity on the social web*, pages 35–40. ACM.
- Marcos Zampieri, Shervin Malmasi, and Mark Dras. 2016. Modeling language change in historical corpora: The case of Portuguese. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*, Paris, France. European Language Resources Association (ELRA).