Political Science Research and Methods (2026), 1–20
doi:10.1017/psrm.2026.10094

ORIGINAL ARTICLE

Measuring inter-party communication: a transformer-based approach

Anna Adendorf[1], Oke Bahnsen[1], Thomas Gschwend[1], Lena Maria Huber[2], Simone Paolo Ponzetto[3], Ines Rehbein[3] and Lukas F. Stoetzer[4]

[1]Department of Political Science, University of Mannheim, Mannheim, Germany; [2]MZES, University of Mannheim, Mannheim, Germany; [3]Data and Web Science Group, University of Mannheim, Mannheim, Germany and [4]Department of Philosophy, Politics and Economics, Witten/Herdecke University, Witten, Germany
Corresponding author: Lena Maria Huber; Email: lena.huber@uni-mannheim.de

(Received 29 November 2023; revised 31 March 2025; accepted 4 May 2025)

Abstract
Inter-party communication is crucial in representative democracies, facilitating information exchange and dialogue among political parties. Despite its importance, research on this topic remains limited due to lacking conceptual clarity and challenges in large-scale measurement. This article offers a comprehensive definition of inter-party communication as public communication by parties about others, with a positive, neutral, or negative stance, focusing on collaboration, policy, or personal issues. To effectively measure this phenomenon, we introduce a novel transformer-based approach capable of automatically classifying large volumes of text. Case studies on coalition signals in Germany and negative campaigning in Austria demonstrate its effectiveness. The study deepens our understanding of party competition, advances methods of automated text classification, and enables new research on political communication.

Keywords: coalition signals; inter-party communication; negative campaigning; transfer learning; transformer models

1. Introduction

Why do parties spend so much time talking about their rivals, partners, or potential allies in public? Effective inter-party communication is essential in representative democracies because it helps the public to monitor and evaluate political parties and make informed electoral choices. Public communication about their opponents can also help parties achieve specific electoral and ideological goals. Despite its importance, research on inter-party communication remains limited, most likely due to conceptual ambiguities and challenges in large-scale measurement.

We define inter-party communication as public communication by political parties about other parties, conveying a positive, neutral, or negative polarity. This communication is purposeful and directed, forming a dynamic network that evolves over time. Parties are connected by their communication signals while the public observes this network, either directly or through the news media.

Elections put parties in competition, which naturally results in the need to publicly negotiate different issues. Most importantly, this includes discussions about coalition formation, policy positions and issue priorities, and the integrity of political elites. Accordingly, our framework integrates three dimensions of inter-party communication that were previously described in isolation, namely communication about collaboration, policy, and personal issues.

This conceptual definition calls for a powerful measurement strategy that can grasp the nuances and context of communication between parties. Therefore, we develop a new transformer-based measurement to automatically classify large amounts of text. Our approach is resource-efficient and more fine-grained, as it is pre-trained on a large corpus of text. It captures context and polarity within sentences and adapts to different types of inter-party communication. We provide a detailed outline of the steps required to implement a transfer learning approach for applications to different types of inter-party communication.[1]

We then showcase our approach with two case studies on coalition signals in German newspaper articles and negative campaigning in Austrian party press releases. Both cases demonstrate the advantages of our method over existing dictionary and supervised approaches. Fine-tuning a pre-trained transformer-based language model gives higher recall and precision. This illustrates how scholars can apply our method when interested in identifying different types of inter-party communication in textual data.

Our study makes two main contributions: a conceptual framework for inter-party communication and an empirical measurement approach using automated text classification.

The conceptualization of inter-party communication allows for subsuming different types of communication under a common umbrella, which up to now have been mostly studied in isolation. Bringing together communication about collaboration, policy, and personal issues, our approach enables a more comprehensive understanding of how parties interact and engage with one another in both adversarial and cooperative contexts, capturing the full spectrum of their strategic behaviors and public messaging. By systematically analyzing different forms of inter-party communication, researchers gain a more nuanced view of the information environment voters are exposed to. A key advantage of this framework is that it enhances our ability to assess the impact of political discourse on public opinion, electoral outcomes, and democratic processes. Furthermore, by establishing a common conceptual language, our framework fosters cross-disciplinary learning, allowing scholars from diverse fields to collaborate more effectively and share methodological insights.

In addition to its theoretical advances, this study contributes to the field of automatic text classification by demonstrating the benefits of fine-tuning transfer-learning models over traditional dictionary-based and supervised machine-learning approaches (Barberá *et al.*, 2021). This improved measurement technique enhances the accuracy and scalability of analyzing large volumes of political text, making it a valuable tool for future studies. Our approach opens new avenues for empirical research by enabling the automated classification of inter-party communication across various textual sources, allowing researchers to investigate critical questions with greater efficiency.

2. Conceptualizing inter-party communication

Public communication between various political parties is part of the larger phenomenon of political communication and represents an important characteristic of democracies. Following conventional definitions, we understand political communication as "purposeful communication about politics" (McNair, 2018, p. 4). Political communication can occur between three key actors: political elites, the media, and the citizens (e.g., Zaller, 1999; McNair, 2018; Dumdum and Bankston, 2022). It is driven by the goal of parties and other political elites to gain public support, the interest of the media in maximizing their audience, and the motivation of citizens to hold political elites accountable (Zaller, 1999). This article focuses on political communication between various political parties, what we call *inter-party communication*.

[1]By transfer learning, we refer to settings where there is actual learning involved, i.e., where the model parameters are updated. This contrasts with transfer settings where a model trained on dataset A is applied to data from a new domain B. We argue that in this setting, no actual learning is involved as the model remains the same (i.e., there is no parameter update).

Despite its importance, research on inter-party communication remains scarce. According to a comprehensive review of research articles published in leading communication and political science journals between 2000 and 2017, only 7.06 percent of the articles examined communication among political elites (Dumdum and Bankston, 2022). Naturally, an even smaller proportion of articles focuses on inter-party communication. The scarcity of empirical research in this field is also reflected by the lack of conceptualizations of inter-party communication (for an exception, see De Nooy and Kleinnijenhuis, 2013).

We define inter-party communication as purposeful communication between political parties. Specifically, following McNair (2018, p. 4), we describe inter-party communication as *all forms of communication undertaken by political parties about other political parties for the purpose of achieving specific objectives*. Each instance involves a purposeful statement made by one party about another, either with a positive, neutral, or negative stance.

Building on the work by De Nooy and Kleinnijenhuis (2013) (see also Song *et al.*, 2019, for a similar conceptualization), we propose a network-centric perspective on inter-party communication. This perspective envisions inter-party communication as a dynamic network (Snijders *et al.*, 2010), consisting of various political parties (*nodes*) and statements about other parties (*ties*) with varying polarities (positive, neutral, or negative). These ties are asymmetric, meaning that reciprocity is not required. The network is dynamic because the ties between the parties change over time.

Figure 1 illustrates such a network of three parties at a given point in time. Here, Party A sends a positive message about Party C, Party B makes a neutral statement about Party A, while Party C sends a negative message about Party B. In this example, there is no reciprocity. This illustration also emphasizes the interdependence of all three key actor groups in political communication, including the media and the citizens. Party communication reaches citizens either directly or via the media, that media coverage influences political parties (e.g., via agenda-setting) and the interests of citizens shape media reporting, as well as the communication between parties (see, e.g., McNair, 2018, Chapter 1).

Existing research on party competition predominantly revolves around three main aspects: issue competition (Green-Pedersen and Mortensen, 2015) and policy positions (Adams, 2012), negative campaigning (Dolezal *et al.*, 2016), and personalization (Pedersen and Rahat, 2021). These areas of research demonstrate how parties communicate strategically about themselves and their political rivals to influence voters, discredit opponents, and position themselves within the political landscape. Our conceptualization integrates these strands and extends beyond them. Specifically, we propose that inter-party communication encompasses three dimensions: collaboration, policy, and personal issues. This framework synthesizes key insights from the literature on party competition and political communication, providing a structured approach to examining the complex interactions between parties, both in adversarial and cooperative contexts. By distinguishing between these three dimensions, we can better analyze the strategic motivations behind party discourse and its implications for political competition, voter information, and coalition dynamics.

The first dimension (collaboration) refers to all communication signals where a party takes a positive, neutral, or negative stance toward another party in terms of their willingness to collaborate or cooperate. This includes statements or actions[2] that indicate support for joint initiatives, alliances, or coalition formation, as well as those that signal opposition or reluctance to work together. Neutral signals indicate ambiguity or non-commitment toward cooperation with another party.

The second dimension (policy issues) reflects whether a party expresses a positive, neutral, or negative attitude toward another party's policy positions, issue emphasis, pledges, or record (Dolezal *et al.*, 2018). This involves examining statements or actions that endorse or critique the policy proposals or ideological stances of another party.

[2]While our focus primarily lies on verbal communication, non-verbal cues, such as applause, gestures, or facial expressions, can also play a significant role for inter-party communication (see, e.g., Imre *et al.*, 2023).

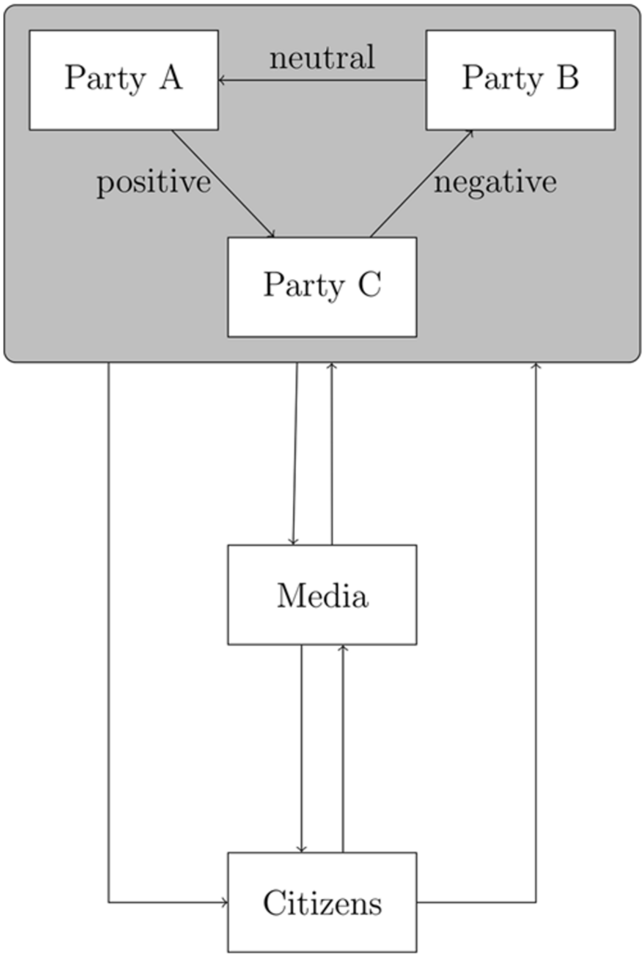

Figure 1. Inter-party communication network in a three-party system.
Note: The arrows within the gray shaded square illustrate purposeful statements with a positive, negative, or neutral stance between parties.

The third dimension (personal issues) refers to inter-party communication about individual representatives of rival parties. It allows for the assessment of whether a party's communication conveys positive or negative stances toward the leaders, candidates, or other personnel of another party. This encompasses statements or actions that praise or criticize the competence, character, or behavior of individuals within the opposing party.

Different types of inter-party communication can have varying effects on voters and other political actors, which explains why some forms are more frequent than others. For instance, while negative statements about a competitor are common, coalition signals (see, e.g., Meffert and Gschwend, 2011; Gschwend *et al.*, 2017) can have severe electoral consequences or change entire campaign strategies. Still, we believe it is valuable to integrate all these aspects into a unified framework because this allows for a more comprehensive analysis of political discourse and its effects, as well as the interconnectedness between these different types of communication.

It follows from this definition that inter-party communication has three inherent characteristics. First, inter-party communication is *purposeful*: It serves parties' primary goal of maximizing votes in elections. Accordingly, statements made by parties about other parties seek to influence the voting calculus of citizens. As citizens learn about party communications mainly through media coverage, parties have an incentive to make newsworthy statements to attract media attention (Strömbäck, 2008).

Second, inter-party communication is *directed*: Each instance of inter-party communication involves a sender party and one or more addressee parties. While reciprocity is not required, empirical evidence suggests that attacks from one party often lead to counterattacks from the other party, following a tit-for-tat strategy (e.g., Dolezal *et al.*, 2016).

Third, inter-party communication is *dynamic*. It is continuously ongoing and changes over time. This is because the logic of the electoral cycle influences public statements by parties about others. For instance, government parties may speak positively about each other throughout the legislative term to promote government unity but become more critical during election campaigns to win votes. Moreover, parties respond to prior attacks or support from other parties, resulting in an evolving communication network (De Nooy and Kleinnijenhuis, 2013).

Due to its dynamic nature (change over time) and directionality (lack of reciprocity), inter-party communication tends to be subtle and diffuse. An important requirement for the empirical measurement of inter-party communication is, therefore, to be able to grasp the nuances and context of communication between parties.

3. Existing approaches to measure inter-party communication

Previous studies on specific types of inter-party communication analyze both mediated (newspaper articles or TV news [Lau and Pomper, 2002; De Nooy and Kleinnijenhuis, 2013]) and unmediated communication channels (campaign ads [Walter and Brug, 2013], press releases [Dolezal *et al.*, 2016], manifestos [Dolezal *et al.*, 2018], or social media [Auter and Fine, 2016]). A common method is manual coding, where human coders classify text by stance and target to study party interactions (see for example Dolezal *et al.*, 2016). Manual coding offers detailed insights and precise information, as researchers can develop specific coding instructions tailored to their research questions.

Yet, there are several drawbacks to manual coding. First, it is time- and resource-intense as multiple coders are required to ensure the reliability of the results. Second, the data obtained through manual coding is tailored to a specific research question, making applications for new analyses and different cases difficult. Hence, to efficiently collect data on inter-party communication, a manual approach is not an ideal solution. This is especially true when studying phenomena that recur in every election, such as negative campaigning or pre-electoral coalition signals. It is, therefore, critical to develop new methods for collecting these data in an automated manner that can be readily applied to new cases.

Dictionaries, such as the one developed by Bowler *et al.* (2022) to detect pre-electoral coalition signals, offer a more scalable alternative. Although this dictionary currently identifies only positive signals between parties, it represents a significant step toward automating the study of inter-party communication as it can be applied to other cases than the one under immediate study. However, dictionaries often fail to capture the nuances and context of communication including tone, sarcasm, or implicit messages. Since the nature of political discourse is highly complex, understanding the underlying meaning and implications of statements demands a deeper level of analysis.

To address these shortcomings, we propose a transformer-based approach that can capture the nuances, context and polarity in party rhetoric more accurately and adapts across domains with minimal training data. This enables large-scale, fine-grained analysis of inter-party communication across different sources and languages (Laurer *et al.*, 2024). Because it can rapidly process vast amounts of data from different sources, this represents an ideal solution for studying inter-party communication across various elections and political landscapes.

4. A transformer-based approach for the classification of inter-party communication

We introduce a transformer-based approach for measuring different types of inter-party communication, building on recent advances in Natural Language Processing (NLP), specifically transfer learning based on pre-trained transformer-based language models (Vaswani *et al.*, 2017; Devlin *et al.*, 2019). These models are already pre-trained on large corpora and therefore capture contextual information more efficiently and require less labeled training data for fine-tuning on new tasks. This is referred to as *transfer learning*, where knowledge learned through training on one task is transferred to solve another task, by using the model parameters of the first model as *a priori* information when learning the parameters of the second model.

While transformer-based models are increasingly used in political science across a range of topics (e.g., Bonikowski *et al.*, 2022; Müller and Proksch, 2024; Widmann and Wich, 2023; Müller and Fujimura, 2025), our contribution focuses on interactions between political actors. We analyze both the content and valence of inter-party communication, distinguishing positive, neutral, and negative stances to capture how parties engage with one another.

For our task, we apply and evaluate a well-known transformer architecture, called BERT (Devlin *et al.*, 2019).[3] Our implementation is based on the Hugging Face Transformers library (Wolf *et al.*, 2020) and PyTorch (Paszke *et al.*, 2017). For more details on the training process, downsampling, and hyperparameter settings, please refer to Appendix A. To enhance the model's performance for our specific task, we add a fine-tuning step, using a small set of annotated data.

We now describe the individual steps of our general approach. Different types of inter-party communication can be modeled similarly. First, we need to identify the target of the communicated message. Second, we need to determine the (positive, negative, or neutral) stance of the message toward the target. To model this, we decompose the problem into two separate tasks and train a classifier for each task. The first step consists of identifying sentences that contain a probable *target* and predicting the target name. In the second step, we learn to predict the *polarity* of the message. Modeling both tasks in one step would result in numerous sparse label combinations of target + polarity that are infeasible to learn for any machine learning model. Breaking down the task into two steps reduces the number of labels and enhances the efficiency of training data usage. Polarity determination in a message or signal should not rely on the specific target, making this approach more effective. The same model architecture is used for both tasks.

We compare our approach to a traditional supervised machine learning system, and we also include a dictionary approach as a baseline for the performance evaluation. Figure 2 outlines the workflow for the dictionary approach (left), the traditional machine learning classifier (middle), and the transfer learning approach (right).

To ensure the validity and applicability of our method in real-world political scenarios, we test the classification for both mediated and unmediated forms of communication. In Case Study 1, we predict coalition signals in German newspaper articles and in Case Study 2, we study negative campaigning in Austrian party press releases. We focus on coalition signals and negative campaigning, as these are the most prominent types of communication between political parties.

5. Case study 1: identifying coalition signals in the news

The first case study applies our transformer-based approach to detect pre-electoral *coalition signals*. Coalition signals are defined as statements about a party's preference toward a possible coalition in which the party itself would be a member (Gschwend *et al.*, 2017). These signals are common in parliamentary democracies, where the formation of coalitions after the election to form a government is typically required. They can be regarded as inter-party communication on collaboration with either a positive, neutral or negative stance polarity.

A coalition signal has three key components: the sender, the addressee, and the statement's polarity. For a statement to be identified as a coalition signal, it is not required that the sender directly speaks about the party's preferences concerning a specific coalition option, but it might also include

[3] We also tested two other transformer models, RoBERTa (Liu *et al.*, 2019) and XLNet (Yang *et al.*, 2019), but obtained higher F1 scores on our data using BERT.

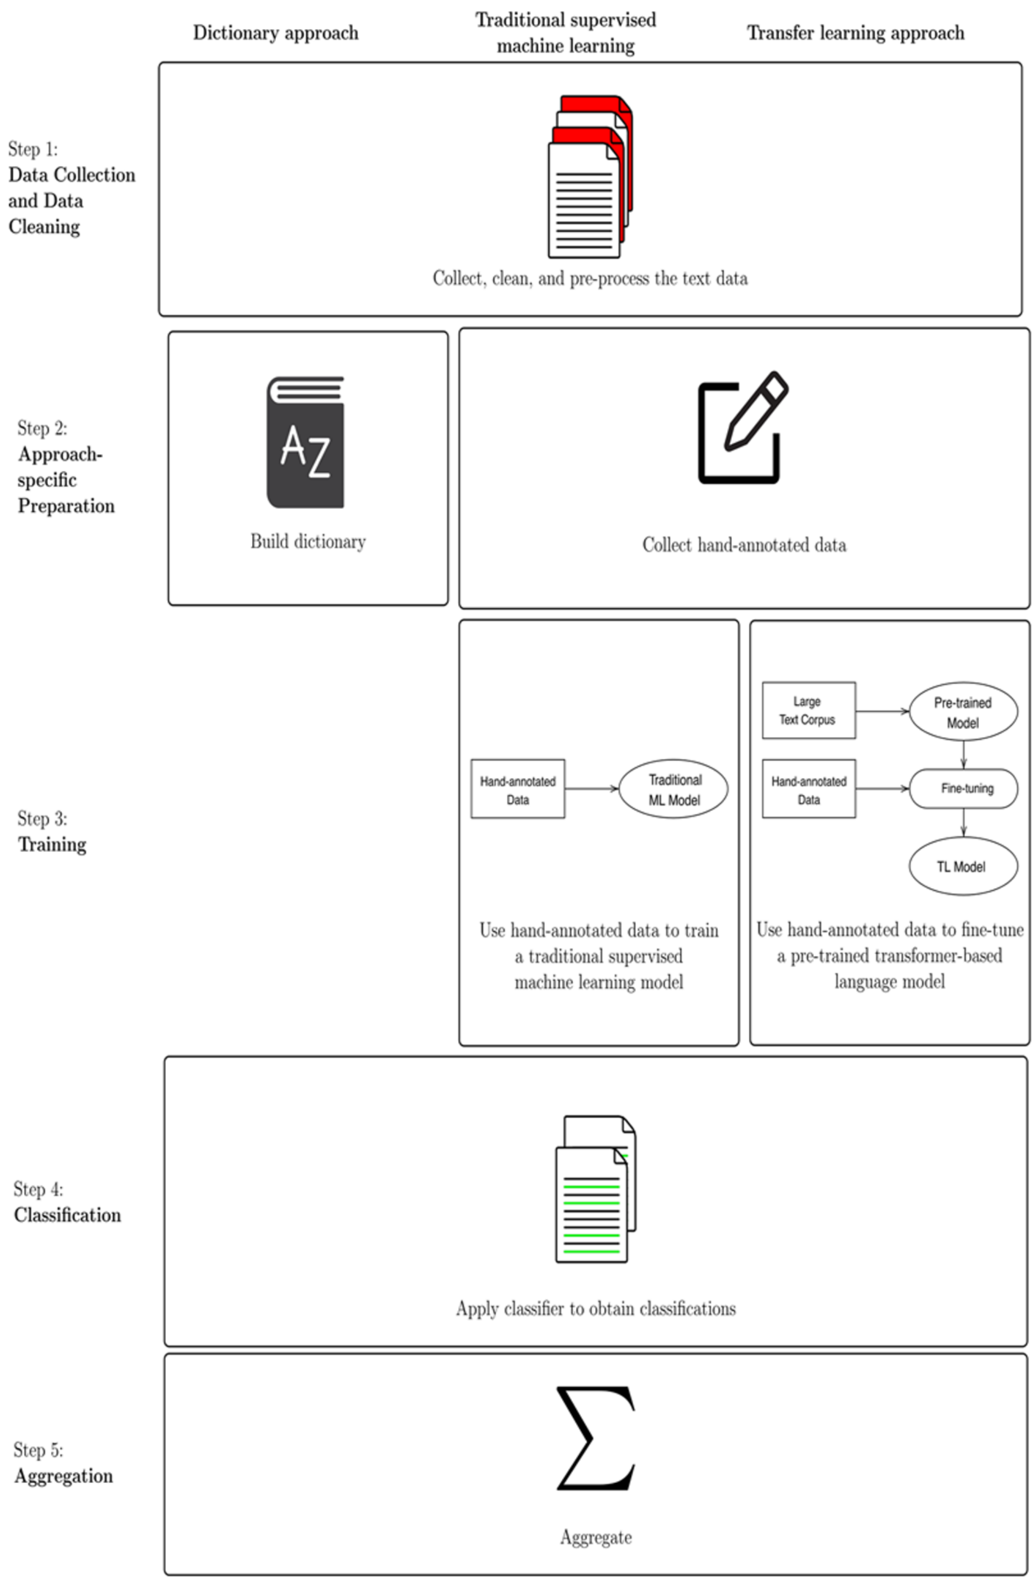

Figure 2. Three approaches to classify elite communication.

https://doi.org/10.1017/psrm.2026.10094 Published online by Cambridge University Press

the citation of (direct and indirect) quotes by a journalist. Sometimes, the addressee might be left unspecified (e.g., when a party rules out a coalition with "any left party"). In addition, we also consider a text segment to be a coalition signal if it is not completely clear whether the identified statement about a potential coalition was originally made by a representative of an involved party or merely by an observer, such as a journalist, expert or another political actor. This includes cases when it is not possible to identify the exact source within the party.

5.1. Data creation and prediction of coalition signals

To identify coalition signals in newspaper articles, we rely on three different approaches: the dictionary approach of Bowler *et al.* (2022), a traditional machine learning approach, and our transfer learning approach.

As the first step in the process, we collect and clean the data (see Step 1 in Figure 2). Specifically, we compile German newspaper articles published prior to the German Bundestag elections between 1998 and 2017. We include all articles from two German daily newspapers[4] that were published within four weeks (28 days) before federal elections (not including the actual election day). Relevant articles containing coalition signals were identified through a keyword search. Each selected article had to include both a reference to a party or coalition term and a mention of a federal or state election.

In the second step of the process, we set up the approach-specific preparations (see Step 2 in Figure 2). For the dictionary approach, we follow Bowler *et al.* (2022) and set up a dictionary with party names and terms indicating cooperation or coalition to identify sentences that might include coalition signals. If a sentence includes a term signaling cooperation *and* a reference to a coalition, it is classified as a coalition signal. To be considered as a coalition reference, the sentence must include at least two mentions of different parties *or* a coalition option that clearly identifies the parties participating in the coalition (e.g., "grand coalition" stands for a coalition between CDU/CSU and SPD or a "Jamaica coalition" refers to a coalition between the CDU/CSU, FDP and Greens).

To create a gold standard for training our classifier, we manually annotated newspaper articles. The annotators first read the whole article and highlight all signals in the article.[5] After this initial step, all highlighted signals in the article are coded in chronological order. For each identified coalition signal, we code the sender(s) and addressee(s) of the signal, so that each can be attributed to one or more specific parties. In cases where this information is not explicitly specified in the text, the instance is coded as *unspecified*. Moreover, the annotators identify and code the polarity of each signal (negative, neutral or positive).

Inter-coder reliability for the annotations yielded a Krippendorff's alpha of 0.595 for the proportion of signals per article across 34 articles read by all three annotators. While acceptable by literature standards, this is not particularly high. This should be kept in mind when evaluating the results of our transfer-based approach. In other words, while machine learning techniques are improving rapidly, even these models are likely to struggle with classifying concepts that trained human annotators find difficult to identify in texts.

We then create a training dataset to learn predicting coalition signals between parties in unseen text. Each training instance encodes a signal between two or more parties. Because the number of coalition options becomes too large when also considering the direction of the coalition signal, we map sender and addressee information to undirected coalition options.[6] As a result, we obtain a set of 26 different undirected coalition options. Overall, we find 755 instances of coalition signals in

[4]These are the *Süddeutsche Zeitung (SZ)* and *Frankfurter Allgemeine Zeitung (FAZ)*, which are two of the most important and influential German daily newspapers.

[5]All three annotators were students of political science who received extensive training and followed the instructions of a detailed codebook.

[6]When encoding the direction of the signal, we end up with 39 different coalition options. Also encoding the polarity of the signal further increases the number of labels to 86, making the task infeasible.

our corpus of newspaper articles. Only 12 out of the 26 different coalition options appear at least 10 times in our data (see Table A.2 for an overview). When we also consider the polarity of the (undirected) coalition, the number of different outcomes (coalition × polarity) increases to 57, rendering the prediction task infeasible.[7]

As described above, we decompose the task of automated coalition signal detection in newspaper articles into two subtasks. We model them as a sequence classification problem where we present the model with a sentence and let it learn, i.e., predict first the coalition option (including "none") and then the polarity of the respective sentence (see Step 3 in Figure 2). As we only have a few instances for training, we use an n-fold cross-validation (e.g., Neunhoeffer and Sternberg, 2019) setting, where we train the model on five of the six elections and test the trained model on the unseen sixth election. We repeat this procedure six times in order to obtain predictions for each of the six elections.

As an additional baseline of comparison, we follow the same cross-validation setup as described above and train a supervised, feature-based machine learning classifier on our data. Specifically, we use a Support Vector Machine (SVM) algorithm, which has been shown to obtain good results for text classification problems.[8] We pre-process our input data by tokenizing the data and removing stop words. Then we create the feature vectors that are the input to our SVM. For feature extraction, we use a "bag of words" approach and weigh the extracted features, based on the term frequency–inverse document frequency (TF-IDS) weighting scheme.

During training, we perform Bayesian optimization over hyperparameters on the training set in a stratified 10-fold setup, using Bayes search. Then we use the best parameter setting for each training fold to predict coalition signals in our unseen test set. In the end, we collect all predictions for the targets and the respective polarity for the six test folds and evaluate them against the annotations (see below for the results of this evaluation).

We then apply both classifiers to the newspaper articles to provide us with sentence-level predictions of coalition signals (see Step 4 in Figure 2). The dictionary method by Bowler *et al.* (2022) does not encode the polarity of coalition signals. In contrast, the transfer learning approach is able to differentiate between positive, negative, and neutral coalition signals. Both methods, however, cannot distinguish between the sender and addressee of a coalition signal.

After obtaining sentence-level classifications of coalition signals, it is necessary for many applications to generate a summary measure for a particular election campaign or text (see Step 5 in Figure 2). One potential aggregation measure of coalition signals was introduced by Bowler *et al.* (2022) and indicates how often a coalition option c is mentioned compared to all other coalition options before a particular election e. This simple measure is defined as the quotient of the number of sentences that contain coalition signals concerning coalition option c, $n_{e,c}$, and the total number of sentences that contain coalition signals (irrespective of the coalition signaled), N_e: $\theta_{e,c} = \frac{n_{e,c}}{N_e}$. As a relative frequency, this measure ranges from 0 to 1. $\theta_{e,c} = 0$ implies that there is not a single coalition signal concerning coalition option c before election e, while $\theta_{e,c} = 1$ means that all coalition signals are about coalition option c. As this measure does not account for the polarity of coalition signals, we refer to $\theta_{e,c}$ as *salience* measure of coalition signals.

5.2. Evaluation

We assess the performance of our transformer-based approach based on the hand-coded German newspaper articles, as outlined earlier. For our annotated dataset, we gather sentence-level predictions from our transfer learning method, as well as the dictionary approach by Bowler *et al.* (2022) and the

[7]We illustrate the difficulty of this task in Table A.3 in the Appendix.

[8]Previous assessments validated that the SVM algorithm consistently outperformed other supervised classification methods in terms of performance for our application.

SVM classifier, and evaluate these predictions against the manual annotations. We use the precision, recall and F1-Score to evaluate the performance. The transformer-based predictions (Transformer) in the third row of Table 1 result from a sixfold cross-validation setup, generating out-of-sample predictions for each of the six test folds. We additionally report standard deviations over five independent runs with different model initializations. Finally, to account for potentially large variance across different initializations, we apply a well-known machine learning technique, i.e., ensemble learning (Dong *et al.*, 2020), where we consider the five models as an ensemble of classifiers and determine the predicted labels by taking the majority vote over the classifier ensemble which we report in the fourth row (TransformerE).

We present the results for four evaluation settings (A, B, C, and D), summarized in Table 1. In setting (A), we compare the three approaches at the sentence level, assessing the correct prediction of coalition signals without considering coalition option or polarity. The results show a higher F1 score for our transformer-based approach (50.4%) than for the dictionary approach (38.8%) and the SVM (44.1%). Although the dictionary approach has a higher precision than the transformer, this clearly comes at the cost of recall (44.5% versus 84.9%).

In settings (B) and (C), we shift the focus to the signal level. However, challenges arise because hand annotations frequently span multiple sentences per signal, while predictions are at the sentence level. By mapping predictions of individual sentences to corresponding signals, we are able to evaluate the performance on the signal level. Setting (B) evaluates predicted coalition options regardless of polarity while setting (C) considers both coalition option and polarity.[9] The F1 score in setting (B) is higher for our transfer learning approach (64.9%) than for the competing dictionary approach (52.9%) where our approach again has a lower precision but a higher recall. Results for the SVM are even below the dictionary baseline (47.9%). In setting (C), the accuracy as measured by the F1 score is again higher for our approach (44.1%) than for the dictionary approach (30.5%) and the SVM (31.3%). On average, the transformer-based approach accurately predicts 350 out of the 755 coalition signals in the data,[10] whereas the dictionary and SVM approaches only detect 185 and 206 signals, respectively. We can conclude that our approach outperforms both baselines.

For the prediction of the ensemble classifier (TransformerE), we merely collect the predictions made by the five transformer-based models and assign each instance the label that has been predicted by the majority of the classifiers. We then evaluate the final labels against the manually assigned labels. We observe improved F1 scores consistently across all our settings.

Setting (D) extends the evaluation beyond the sentence- and signal-level by aggregating coalition signals to the election-level. We compute the salience values for both the machine learning approaches and the dictionary approach and measure the difference between the predicted and the true salience values using the root-mean-square error (RMSE). The RMSE for our transfer learning approach is lower than that for the dictionary approach, showing that the salience values for election-coalition combinations predicted by our approach are slightly closer to the true values than the salience values by the dictionary approach. The error for the SVM, however, is notably higher, despite obtaining a slightly higher F1 in Setting C. This reflects the lower precision and F1 of the model in setting B compared to the dictionary approach, as the salience measure does not take into account the polarity of the signal. We also test for statistical significance for all models, using the McNemar test (for details see Appendix A) and find that the difference in model performance for the dictionary approach and the SVM is not statistically significant while the increase in performance between the transformer models and the dictionary is strongly significant ($\alpha = 0.0001$). We also test for differences between the individual transformer models and find no significant differences, with p-values ranging between 0.1 and 1.0.

[9]Arguably, setting (B) is a fairer comparison than setting (C) as the dictionary approach does not model the polarity of coalition signals.

[10]The number of correct predictions is averaged over the 5 runs with different initializations.

https://doi.org/10.1017/psrm.2026.10094 Published online by Cambridge University Press

Table 1. Results for the prediction of coalition signals in newspaper articles

	(A) Signal (yes/no)			(B) Coalition			(C) Coalition, polarity			(D) Aggregation
	Prec	Rec	F1	Prec	Rec	F1	Prec	Rec	F1	RMSE
Dictionary	34.3	44.5	38.8	70.1	42.5	52.9	40.4	24.5	30.5	0.0014
SVM	37.1	54.4	44.1	56.9	41.3	47.9	36.6	27.3	31.3	0.0019
Transformer	32.5 ± 0.3	84.9 ± 0.6	50.4 ± 0.4	63.2 ± 0.4	66.7 ± 0.7	64.9 ± 0.5	42.1 ± 0.5	46.4 ± 0.5	44.1 ± 0.5	0.0008
TransformerE	34.1	84.7	48.6	63.6	66.5	65.0	43.3	47.3	45.2	0.0009

Note: (A) Prediction of coalition signals on the sentence-level (signal: yes/no); (B) prediction of the coalition option on the signal-level; (C) prediction of the coalition option *and* polarity on the signal-level; and (D) aggregation measure $\theta_{e,c}$. For the Transformer approach, we additionally report standard deviations over 5 independent runs with different model initializations. TransformerE reports results for an ensemble classifier with majority vote, based on the predictions of the five Transformer models.

In sum, we find that our transformer-based approach outperforms both the dictionary approach of Bowler *et al.* (2022) and the SVM classifier on each of the evaluation settings described above, as indicated by the F1 scores and the RMSE.

5.3. Results

Figure 3 illustrates the network of coalition signals in newspaper articles for German national elections between 1998 and 2017. To provide a clear overview of both the parties involved in coalition signals and the overall balance of positive and negative signals, we aggregate our results for each election at the dyadic level – that is, between each pair of parties.[11]

In this network, parties are represented as nodes, with node size indicating the number of coalition signals involving a party as a sender or addressee. This reflects how prominently a party features in pre-election discussions about potential coalition formations. The width of the edges represents the relative strength of inter-party communication, measured as the balance between positive and negative coalition signals. Following Lowe *et al.* (2011), who applied a similar approach to measuring left-right positions in party manifestos, we quantify this balance as $\log \frac{n_{e,c}^{pos}+0.5}{n_{e,c}^{neg}+0.5}$ where $n_{e,c}^{pos}$ is the number of sentences containing a *positive* coalition signal for a coalition option c before election e, and $n_{e,c}^{neg}$ is the number of sentences containing a *negative* coalition signal for the same coalition. An overall positive signal is represented in blue, while a negative signal appears in red.

The figure clearly shows that the identifiability of alternative coalition options before the 1998 election was stronger than at any other point in recent history. Two distinct coalitions were competing: a right-leaning CDU/FDP alliance and a left-leaning SPD/Greens coalition. This gave voters a clear sense of how their party vote could also influence the composition of the next government. While a right-leaning coalition remained a viable option across elections – albeit with varying degrees of strength – the political left struggled to present a similarly functional two-party bloc. Across most elections, the signal for a left-leaning coalition was not always positive, and even when it was, it often lacked strength. This reflects the perception that an SPD/Greens coalition was hardly ever popular enough to secure a parliamentary majority. In both 2013 and 2017, the SPD and the Left Party appeared to emphasize their cooperative relationship in hopes of being seen as a viable governing option together with the Greens. However, the Greens never seemed supportive of this alternative left-wing coalition, as indicated by the weak red or absent edges between them and the Left Party. Additionally, the SPD was also never able to rely on positive signals suggesting a potential majority coalition that included the FDP, such as a so-called "traffic light" coalition of SPD, FDP, and Greens.

6. Case study 2: identifying negative campaigning in party press releases

The second case study evaluates our method for detecting negative campaigning in Austrian party press releases during national election campaigns. Negative campaigning is defined as communication that attacks another party's policy proposals, its past record, or personnel (Lau and Rovner, 2009, p. 286). Hence, negative campaigning can be thought of as inter-party communication with a negative polarity.

6.1. Data creation and prediction of negative campaigning

To study negative campaigning in the context of Austrian elections, we need to identify whether a press release addresses a specific party and whether the author's stance toward this party is positive, negative or neutral. Again, we compare the results for our proposed transformer-based approach to a traditional feature-based SVM baseline and a simple dictionary approach, as described below.

[11]For signals involving more than two parties, such as a "Jamaica coalition" (CDU/FDP/Greens), we increase the counts for each party-dyad and the respective polarity separately.

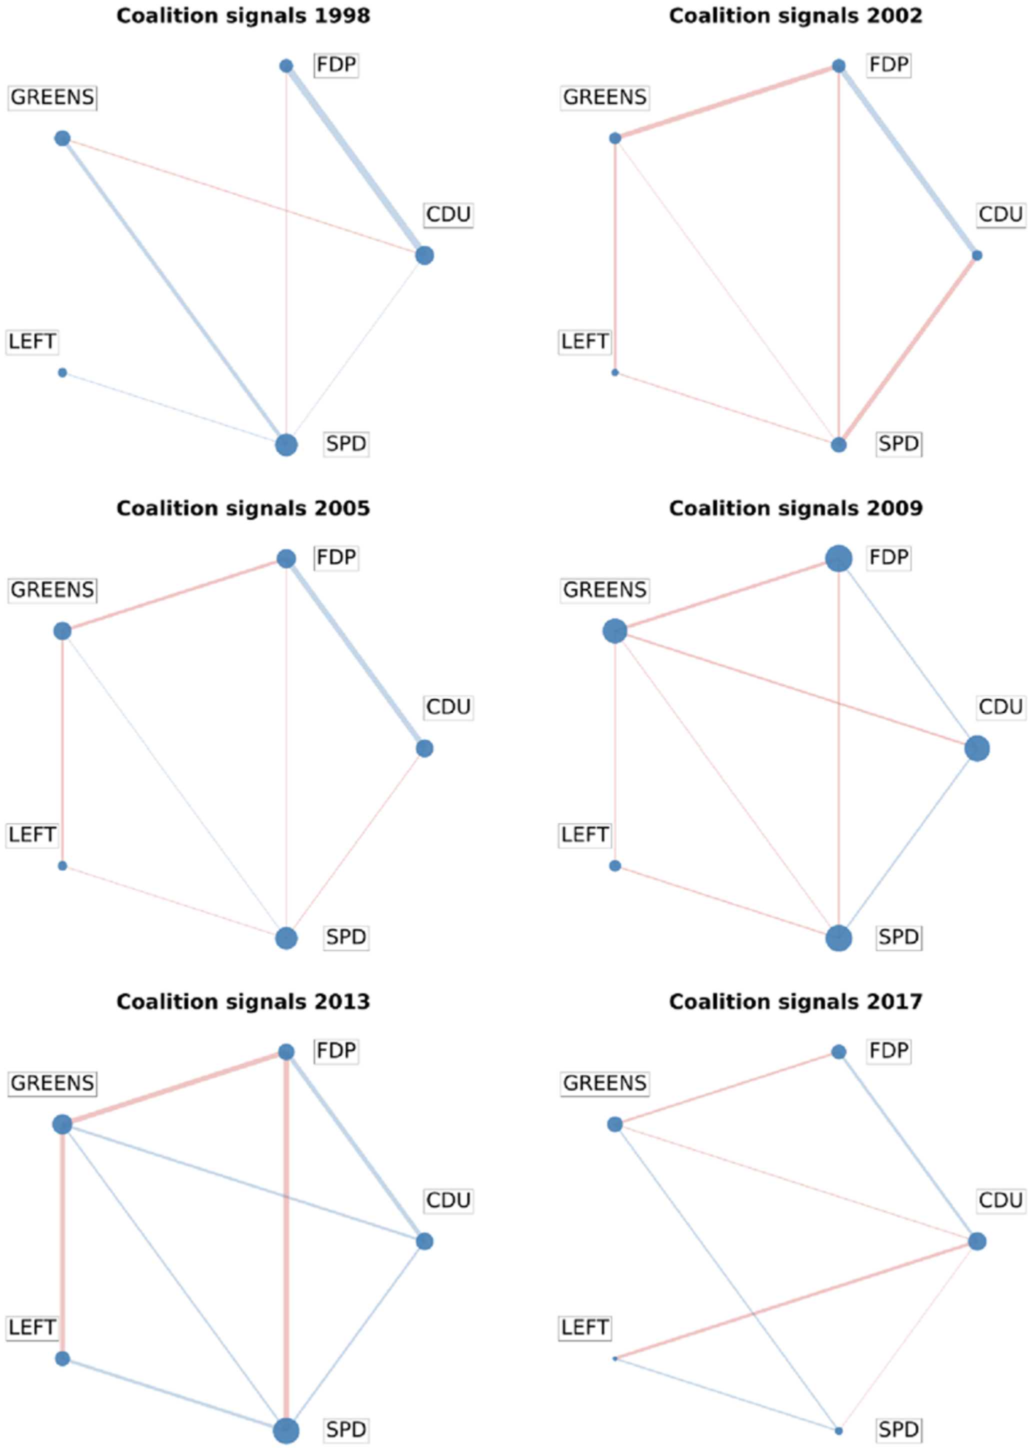

Figure 3. Network of coalition signals between parties.
Note: Blue edges indicate positive signals, while red edges indicate negative signals. The node size indicates the signal frequency, while the edge width shows the signal strength.

We use a data set of party press releases provided by the Austrian National Election Study (AUTNES; Müller *et al.*, 2021), which includes press releases issued by Austrian parties within a time frame of six weeks before the national elections between 2002 and 2017. The data contain the title, subtitle and the content of the press release, together with additional information about the date of publication and the channel used for transmission.

After cleaning the data (see Step 1 in Figure 2), human annotators coded the author of the release, their party affiliation, the issues or actors discussed in the release and the author's stance toward this issue or actor (support, rejection, neutral) (Step 2 in Figure 2). The data cover nearly 10,000 press releases (see Table B.1), coded for 425 different issues or actors, including 12 political parties. We focus on those press releases that target the four major parties that competed in all elections during the respective time frame (SPÖ, ÖVP, FPÖ, Greens) and assign all press releases that are not about one of these parties to the OTHER class. As a result, we end up with a set of five target labels that we want to predict.

Once we know the target, we want to predict whether a press release conveys a positive, negative, or neutral stance toward this target. The prediction of stance, however, is only meaningful for press releases with a target, i.e., that include (positive or negative) statements about a rival party. We remove all press releases without a target (OTHER class) from the training and validation set, assuming that there is no useful information for the classifier to learn. For the prediction, however, we keep these instances in the test set and include them in the evaluation, mapping their label to the negative class. This means that the results for stance prediction are unrealistically low, which is not a problem for us as we are mostly interested in the detection of negative campaigning. For the final evaluation, we rely on the predictions of the target classifier to identify press releases that belong to the OTHER class and only use the stance predictions for the remaining instances.

For the dictionary approach, we create a simple dictionary with party names to identify relevant targets in the press releases (for details, see Table B.2 in the Appendix). To predict the author's stance toward the target, we use SentiMerge (Emerson and Declerck, 2014). We also lowercase the text in the title and subtitle of the press releases and remove stopwords.

Our setup for training follows the approach described above (see Step 3 in Figure 2). Again, we decompose the task into two subtasks, (1) predicting the target (SPÖ, ÖVP, FPÖ, Greens, OTHER) and (2) the author's stance toward the target. We use the same model architectures as before and train a traditional feature-based SVM classifier and a transformer-based transfer learning model (see Appendix B for details on feature extraction and training). As the data includes press releases for five elections, we run a fivefold cross-validation where the data for each election year is once used as the out-of-sample test set.

In the classification step (see Step 4 in Figure 2), we apply our trained classifiers to the test set, following the fivefold cross-validation regime. For the dictionary baseline, we use the party name dictionary to look for relevant mentions in the title or subtitle of each press release and label it as the target. If we find more than one party name in the release, we greedily assign the target label to the first candidate.[12] If no party name is found, we label this instance as OTHER.

For polarity prediction, we iterate over the SentiMerge dictionary and look for entries that occur in the press release. We then sum up the (positive and negative) sentiment scores for those entries and normalize the result by the number of dictionary terms found in the text. We count each dictionary term only once even if it occurs multiple times in the text. As a result, we obtain a score for each press release, based on the sum of all sentiment terms in the release. To predict the author's stance toward the target, we determine a threshold for positive and negative polarity. Therefore, we use a bootstrapping process where we repeatedly draw 1,000 data samples with replacement, which gives us a distribution of the sample means. We then take the 2.5 and 97.5 percentile of the distribution as

[12]We also ran experiments where we included the content of the press release, however, results on the development set decreased for this setting, due to a higher number of false positives.

Table 2. Micro-F1 scores for negative campaigning in Austrian press releases

	2002	2006	2008	2013	2017	Total
Target						
Dictionary	60.2	63.4	63.0	55.4	56.2	59.6
SVM	55.9	57.2	64.1	59.4	65.2	60.4
Transformer	80.3 ± 0.87	77.5 ± 0.72	87.2 ± 0.09	81.4 ± 0.92	85.6 ± 0.33	82.4
TransformerE	81.2	77.8	87.9	82.5	86.1	83.0
Target + Stance						
Dictionary	40.8	38.1	38.3	36.0	33.8	37.4
SVM	51.1	50.9	47.1	46.8	42.4	47.7
Transformer	77.5 ± 0.78	75.7 ± 0.69	84.9 ± 0.19	79.3 ± 0.85	84.6 ± 0.41	80.6
TransformerE	79.8	76.1	85.7	80.3	85.4	81.2

Note: Results for the transfer learning approach are averaged over five independent runs with different initializations (± reports standard deviation). TransformerE reports results for an ensemble classifier with majority vote.

the confidence interval and label every press release with a polarity score below the 2.5 percentile as negative and with a score higher than the 97.5 percentile as positive.

6.2. Evaluation

Table 2 summarizes the results for target prediction for the dictionary baseline (1), the SVM (2), and the transformer-based approach (3). The last column presents the averaged F1 score over the five folds for each election year. The dictionary approach outperforms the SVM in the first two elections (2002, 2006) but results decrease over time. On average, the two baselines obtain results in the same range (around 60% F1). For the transformer-based approach, we report averaged results over five different initializations together with the standard deviation (Table 2, Transformer) over those runs. Our method consistently outperforms both baselines on each election, with an average F1 score of over 80%.

Again, given the high standard deviation between models from different initializations, more robust results might be obtained by aggregating the predictions over the classifier ensemble (TransformerE). This yields an improvement over the individual models for each of the elections, with an average F1 of 83%.

Results for predicting both target and stance are shown in Table 2, rows (5)–(8).[13] To obtain the final labels, we merge the target and stance predictions into one atomic label and evaluate these labels against the manually coded classes. The dictionary approach performs poorly, with an F1 in the range of 34–40%. The SVM yields slightly better results (42–51% micro-F1). Again, the transformer-based approach outperforms both baselines on each individual election year, showing that the context-sensitive semantic representations are superior to symbolic, word-count-based representations (i.e., the TF-IDF features used by the SVM). Combining the predictions of the classifier ensemble achieves best results and improves the averaged results from 80.6% to 81.2% micro-F1, thus increasing the robustness of our prediction model.

6.3. Results

Figure 4 visualizes the network of negative campaigning in Austria from 2002 to 2017. Again, the node size reflects how often a party was involved in negative messaging, while edge thickness reflects intensity of negative signals. We can see that instances of negative campaigning most frequently

[13]We cannot report results for stance prediction alone as we removed press releases without a target from the training data and therefore need the combined labels to identify the OTHER class (press releases with no target).

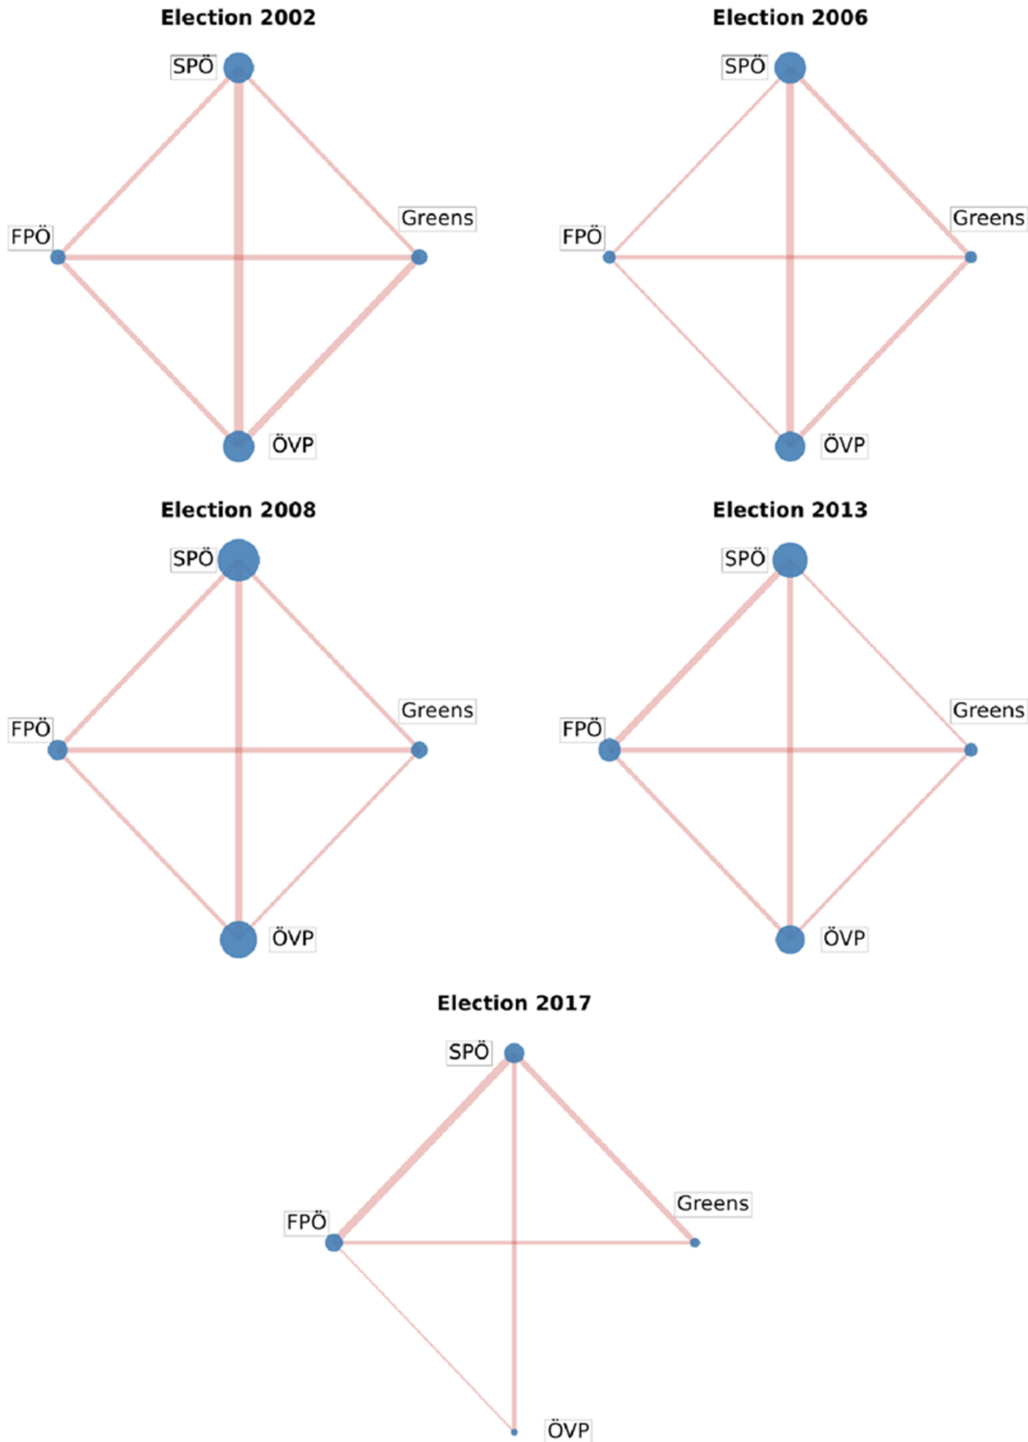

Figure 4. Network for negative campaigning between parties.
Note: Red edges indicate negative signals, the node size indicates signal frequency, and edge width shows signal strength.

occurred between the ÖVP-SPÖ and FPÖ-SPÖ dyads. While ÖVP-SPÖ consistently exhibited the highest number of negative signals (which, however, sharply decline in 2017), FPÖ-SPÖ negativity increased notably from 2008. Other dyads, such as Greens-SPÖ, Greens-FPÖ, and ÖVP-Greens, show fewer instances of negative campaigning, while ÖVP-FPÖ showed moderate levels.

These patterns align with shifting coalition politics and inter-party rivalry in Austria during this period. After the 2002 election, an ÖVP-FPÖ coalition continued, but the FPÖ's internal struggles and declining popularity likely weakened its signaling activity, reflected in the relatively low levels of negative signals involving the FPÖ in the early years. A sharp increase in negativity between ÖVP and SPÖ coincides with a particularly contentious grand coalition (ÖVP-SPÖ) formed after the 2006 election. This coalition was marked by internal friction, public dissatisfaction, and early elections in 2008, all of which likely fueled more aggressive public messaging and signaling. The spike in FPÖ-SPÖ negative signals in 2013 reflect the FPÖ's intensified attacks on establishment parties during this time, as it positioned itself as a strong opposition force ahead of the election. By 2017, negative signaling between ÖVP and FPÖ declined as they were preparing to govern together, reflecting the political realignment that led to the ÖVP-FPÖ coalition under Sebastian Kurz.

7. Discussion

This study explores the understudied yet essential role of inter-party communication in representative democracies. Despite its importance for the facilitation of dialogue, information exchange, and compromise between political parties, research on inter-party communication has remained relatively scarce, partly due to the lack of a comprehensive conceptualization and the challenges associated with large-scale measurement.

Addressing these gaps, we offer a unified conceptual framework that captures how parties engage with one another through public messaging across three dimensions: collaboration, policy, and personal issues. By defining inter-party communication as purposeful, directed, and dynamic, we provide a structure for analyzing party interactions beyond isolated forms of discourse. Thereby we bring together different strands of the literature that rarely speak to one another. This enriched conceptual framework deepens our understanding of interactions between parties and their evolution over time.

Our second contribution is methodological. We introduce a transformer-based approach that significantly improves the large-scale classification of inter-party communication. By leveraging transfer learning, our model captures context and polarity more accurately than conventional methods such as manual coding and dictionary-based techniques. Our method is particularly valuable as it enables the classification of intricate concepts without demanding an excessive volume of training data (see Figure A.2 in the Appendix).

The effectiveness of our approach was demonstrated with two case studies on coalition signals in newspaper articles and negative campaigning in party press releases. The two applications represent particularly frequent and important types of elite communication and directly relate to our conceptualization of inter-party communication. The results show that our approach offers a powerful tool for researchers to analyze textual data on inter-party communication from mediated as well as unmediated communication channels.

Our study has important implications for the study of political communication and beyond. By providing a unified framework for diverse forms of communication, we can now subsume various types of communication that were previously studied in isolation under a common umbrella. This conceptual clarity allows for a more accurate assessment of parties' communication dynamics and enables us to uncover the broader patterns and strategies that drive party competition, collaboration, and negotiation dynamics. Another advantage of a comprehensive conceptualization is that it can promote learning from other subject areas by creating a common language. This will make it easier to develop new standardized measurement approaches, fostering research collaboration across different fields, and potentially reducing costs associated with human annotation in future studies. This reduction in costs and effort can make large-scale studies more feasible and accessible.

In addition, the versatility of our measurement approach has broader implications for the field of automatic text classification. The transfer learning approach for inter-party communication can be applied to diverse data sources from various communication platforms, including press releases, social media postings, speeches, and others. The method can be extended to explore a range of different research questions, such as the evolution of party strategies over time, and the cross-national variations in party communication tactics. More importantly, it could lead to a better understanding of the impact of inter-party communication on voting behavior, public opinion, and political polarization.

While our study bridges a critical gap in political communication research by enhancing our understanding of inter-party communication in representative democracies, it is also limited in several ways: One significant limitation of our current study is its focus on one language, which restricts the generalizability of our findings to other contexts. A promising direction for future research involves exploring the transferability of our approach to other languages which could provide a more comprehensive understanding of inter-party communication in a multilingual setting. Another area for potential enhancement lies in the more distinct analysis of policy issues and personal issues, both with a positive and negative stance. Our current framework integrates these dimensions, but our analysis is constrained by the format of our training data and does not allow for a detailed differentiation between positive and negative stances on policy and personal issues at a more nuanced level. The same applies to the differentiation between sender and addressee party, which is important from a theoretical point of view, but which we were unable to implement due to data sparsity. Future research could address this limitation by developing a more refined dataset that categorizes statements more explicitly along these lines.

Supplementary material. The supplementary material for this article can be found at https://doi.org/10.1017/psrm.2026.10094. To obtain replication material for this article, please visit https://doi.org/10.7910/DVN/W5LGFK.

Data availability statement. The data and replication code supporting the results and analyses presented in the paper are available in the PSRM Dataverse (https://dataverse.harvard.edu/dataverse/PSRM). The AUTNES data is available from the Austrian Social Science Data Archive https://doi.org/10.11587/25P2WR.

Acknowledgements. We like to thank Frank Brettschneider, Stefan Müller, and the three anonymous reviewers for their helpful comments. Support for this research was provided by the German Research Foundation (DFG) under grants STO 1166/1-1 (Stoetzer) and GS 17/8-1 (Gschwend).

References

Adams J (2012) Causes and electoral consequences of party policy shifts in multiparty elections: theoretical results and empirical evidence. *Annual Review of Political Science* **15**(1), 401–419.

Auter ZJ and Fine JA (2016) Negative campaigning in the social media age: attack advertising on facebook. *Political Behavior* **38**(4), 999–1020.

Barberá P, Boydstun AE, Linn S, McMahon R and Nagler J (2021) Automated text classification of news articles: a practical guide. *Political Analysis* **29**(1), 19–42.

Bonikowski B, Luo Y and Stuhler O (2022) Politics as usual? measuring populism, nationalism, and authoritarianism in U.S. presidential campaigns (1952–2020) with neural language models. *Sociological Methods & Research* **51**(4), 1721–1787.

Bowler S, McElroy G and Müller S (2022) Voter expectations of government formation in coalition systems: the importance of the information context. *European Journal of Political Research* **61**(1), 111–133.

De Nooy W and Kleinnijenhuis J (2013) Polarization in the media during an election campaign: a dynamic network model predicting support and attack among political actors. *Political Communication* **30**(1), 117–138.

Devlin J, Chang M-W, Lee K and Toutanova K (2019) BERT: pre-training of deep bidirectional transformers for language understanding. *Proceedings of NAACL-HLT 2019* **1**, 4171–4186.

Dolezal M, Ennser-Jedenastik L and Müller WC (2016) Negative campaigning and the logic of retaliation in multiparty competition. *The International Journal of Press/Politics* **21**(2), 253–272.

Dolezal M, Ennser-Jedenastik L, Müller WC, Praprotnik K and Winkler AK (2018) Beyond salience and position taking: how political parties communicate through their manifestos. *Party Politics* **24**(3), 240–252.

Dong X, Yu Z, Cao W, Shi Y and Ma Q (2020) A survey on ensemble learning. *Frontiers of Computer Science* **14**(2), 241–258.

Dumdum OO and Bankston L (2022) The interplay of actors in political communication: the state of the subfield. *Political Communication* **39**(2), 266–279.

Emerson G and Declerck T. 2014. SentiMerge: combining sentiment lexicons in a Bayesian framework. In Proceedings of Workshop on Lexical and Grammatical Resources for Language Processing. Dublin, Ireland: Association for Computational Linguistics and Dublin City University. 30–38.

Green-Pedersen C and Mortensen PB (2015) Avoidance and Engagement: Issue Competition In Multiparty Systems. *Political Studies* **63**(4), 747–764.

Gschwend T, Meffert MF and Stoetzer LF (2017) Weighting parties and coalitions: how coalition signals influence voting behavior. *The Journal of Politics* **79**(2), 642–655.

Imre M, Ecker A, Meyer TM and Müller WC (2023) Coalition mood in European parliamentary democracies. *British Journal of Political Science* **53**(1), 104–121.

Lau RR and Pomper GM (2002) Effectiveness of negative campaigning in U.S. senate elections. *American Journal of Political Science* **46**(1), 47.

Lau RR and Rovner IB (2009) Negative campaigning. *Annual Review of Political Science* **12**(1), 285–306.

Laurer M, Van Atteveldt W, Casas A and Welbers K (2024) Less annotating, more classifying: addressing the data scarcity issue of supervised machine learning with deep transfer learning and BERT-NLI. *Political Analysis* **32**(1), 84–100.

Liu Y, Ott M, Goyal N, Du J, Joshi M, Chen D, Levy O, Lewis M, Zettlemoyer L and Stoyanov V. 2019. "RoBERTa: A Robustly Optimized BERT Pretraining Approach".

Lowe W, Benoit K, Mikhaylov S and Laver M (2011) Scaling policy preferences from coded political texts. *Legislative Studies Quarterly* **36**(1), 123–155.

McNair B (2018) In *An Introduction to Political Communication*, 6th Edn. London; New York: Routledge.

Meffert MF and Gschwend T (2011) Polls, coalition signals and strategic voting: An experimental investigation of perceptions and effects: polls, coalition signals and strategic voting. *European Journal of Political Research* **50**(5), 636–667.

Müller S and Fujimura N (2025) Campaign communication and legislative leadership. *Political Science Research and Methods* **13**(3), 545–566.

Müller S and Proksch S-O (2024) Nostalgia in European Party Politics: A Text-Based Measurement Approach. *British Journal of Political Science* **54**(3), 993–1005.

Müller WC, Bodlos A, Dolezal M, Eder N, Jedenastik LE, Gahn C, Graf E, Haselmayer M, Haudum T, Huber LM, Kaltenegger M, Meyer TM, Praprotnik K, Reidinger V and Winkler AK (2021) "AUTNES content analysis of party press releases: cumulative file (SUF edition)."

Neunhoeffer M and Sternberg S (2019) How cross-validation can go wrong and what to do about it. *Political Analysis* **27**(1), 101–106.

Paszke A, Gross S, Chintala S, Chanan G, Yang E, DeVito Z, Lin Z, Desmaison A, Antiga L and Lerer A. 2017. "Automatic differentiation in pytorch." NeurIPS Autodiff Workshop.

Pedersen HH and Rahat G (2021) Political personalization and personalized politics within and beyond the behavioural arena. *Party Politics* **27**(2), 211–219.

Snijders TAB, Van De Bunt GG and Steglich CEG (2010) Introduction to stochastic actor-based models for network dynamics. *Social Networks* **32**(1), 44–60.

Song H, Nyhuis D and Boomgaarden H (2019) A network model of negative campaigning: the structure and determinants of negative campaigning in multiparty systems. *Communication Research* **46**(2), 273–294.

Strömbäck J (2008) Four phases of mediatization: an analysis of the mediatization of politics. *The International Journal of Press/Politics* **13**(3), 228–246.

Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser L and Polosukhin I (2017) Attention is all you need. In Guyon I, Von Luxburg U, Bengio S, Wallach H, Fergus R, Vishwanathan S and Garnett R (eds), *Advances in Neural Information Processing Systems*, Vol. 30. Red Hook, NY, USA: Curran Associates, Inc 5998–6008.

Walter AS and Brug WVD (2013) When the gloves come off: Inter- party variation in negative campaigning in Dutch elections, 1981–2010. *Acta Politica* **48**(4), 367–388.

Widmann T and Wich M (2023) Creating and comparing dictionary, word embedding, and transformer-based models to measure discrete emotions in German political text. *Political Analysis* **31**(4), 626–641.

Wolf T, Debut L, Sanh V, Chaumond J, Delangue C, Moi A, Cistac P, Rault T, Louf R, Funtowicz M, Davison J, Shleifer S, Von Platen P, Clara M, Jernite Y, Plu J, Canwen X, Scao TL, Gugger S, Drame M, Lhoest Q and Rush A. 2020. Transformers: State-of-the-Art Natural Language Processing. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations. Online: Association for Computational Linguistics 38–45.

Yang Z, Dai Z, Yang Y, Carbonell JG, Salakhutdinov R and Le QV (2019) XLNet: Generalized autoregressive pretraining for language under- standing. In Wallach HM, Larochelle H, Beygelzimer A, D'Alche-Buc F, Fox EB and Garnett R (eds), *Advances in neural information processing systems 32: Annual conference on neural information processing systems 2019, NeurIPS 2019, December 8-14, 2019*, Red Hook, NY, USA: Curran Associates, Inc. pp. 5754–5764.

Zaller J (1999) *A Theory of Media Politics: How the Interests of Politicians, Journalists, and Citizens Shape the News*. Chicago, IL, USA: University of Chicago Press.

Cite this article: Adendorf A, Bahnsen O, Gschwend T, Huber LM, Paolo Ponzetto S, Rehbein I and Stoetzer LF (2026) Measuring inter-party communication: a transformer-based approach. *Political Science Research and Methods*, 1–20. https://doi.org/10.1017/psrm.2026.10094