

A (More) Realistic Evaluation Setup for Generalisation of Community Models on Malicious Content Detection

Ivo Verhoeven[†], Pushkar Mishra[‡], Rahel Beloch[†]

Helen Yannakoudakis[§], Ekaterina Shutova[†]

[†] ILLC, University of Amsterdam, The Netherlands

[‡] AI at Meta, London, United Kingdom

[§] Dept. of Informatics, King’s College London, United Kingdom

{name.lastname}@uva.nl, pushkarmishra@meta.com, helen.yannakoudakis@kcl.ac.uk

Abstract

Community models for malicious content detection, which take into account the context from a social graph alongside the content itself, have shown remarkable performance on benchmark datasets. Yet, misinformation and hate speech continue to propagate on social media networks. This mismatch can be partially attributed to the limitations of current evaluation setups that neglect the rapid evolution of online content and the underlying social graph. In this paper, we propose a novel evaluation setup for model generalisation based on our few-shot subgraph sampling approach. This setup tests for generalisation through few labelled examples in local explorations of a larger graph, emulating more realistic application settings. We show this to be a challenging *inductive* setup, wherein strong performance on the training graph is not indicative of performance on unseen tasks, domains, or graph structures. Lastly, we show that graph meta-learners trained with our proposed few-shot subgraph sampling outperform standard community models in the inductive setup. We make our code publicly available.¹

1 Introduction

The combination of connectivity and anonymity offered by social media inadvertently also provides the perfect channel for wide-spread dissemination of malicious content (Allcott and Gentzkow, 2017; Müller and Schwarz, 2017; European Commission, 2018; Derron, 2021). By *malicious content*, we consider any form of content detrimental to society, and focus on two specific forms: misinformation and hate speech. Mitigating the effect of malicious content requires content moderation, but this is a labour-intensive process that exacts an immense psychological toll on moderators (Vosoughi et al., 2018; Wiessner, 2021). Consequently, automated detection of malicious content has seen increased

¹Our anonymised code-base is available at: <https://github.com/rahelbeloch/meta-learning-gnns>

academic interest (Ruffo et al., 2023) and industry adoption.

Community models for malicious content detection are models that operate on *social graphs*, i.e., graphs of content and users. They 1) classify content nodes as malicious or not, 2) incorporate information from interacting users in the graph when doing so, and 3) leverage emergent network properties like homophily to boost detection performance (Ma et al., 2021; Hussain et al., 2021). For community modelling on large-scale heterogeneous online communities, Graph Neural Networks (GNNs) are the architecture of choice (Phan et al., 2023).

While community models for malicious content detection perform very well on benchmark datasets (Mishra et al., 2019; Gong et al., 2023), social media platforms continue to grapple with such content. Alharbi et al. (2021) find that high accuracy in malicious content detection is not indicative of trustworthiness in general, as predictions often rely on dataset-specific features. Models also become outdated quickly as online content and communities evolve (Monti et al., 2019). Bozarth and Budak (2020) find detection models to be brittle to changes in domain or publication date, a finding that Nielsen and McConville (2022) corroborate specifically for community models. Finally, Phan et al. (2023) conclude that “we [have] no graph benchmark data for fake news detection in the graph learning community” (p. 22), making any claims of state-of-the-art performance difficult to verify.

Evidently, there exists a mismatch in the performance of community models on research datasets and in more realistic application settings. Research datasets are static; they capture a view of the social graph weeks or months after relevant content has been introduced and spread. Current evaluation practices designed on static graphs are effectively *transductive* (Song et al., 2021b), i.e., they implicitly assume that no new content or users will be

introduced into the social graph, which leads to performance scores that obscure the discussed deficiencies.

In *realistic* settings, new users and content nodes are constantly added to the social graph, and the topic or style of malicious content often radically changes. This is a property inherent to online content and communities (Adamic et al., 2016; Guo et al., 2021). Thus, a successful community model should be able to rapidly adapt to domain shifts in content. Since labelling is prohibitively expensive, adaptation should occur in a few-shot manner. Furthermore, the community of interacting users also evolves. Initially, only a few users take note of some content, but as it gains traction, more and more users interact. To prevent harm from malicious content, detection must occur before wide-spread dissemination. This requires adaptation from a limited exploration of the social graph. Therefore, *inductive* evaluation is needed.

In this paper, we seek to more realistically assess the generalisation capabilities of community models for malicious content detection, specifically making the following contributions:

1. We design a novel evaluation setup based on a few-shot subgraph sampling procedure that test inductive generalisation. The subgraphs are local, contain limited context, and have only a few labels.
2. We test a state-of-the-art community model under this novel evaluation setup on unseen graphs, domains, and tasks. We find them lacking the capacity to generalise.
3. We show that graph meta-learners trained with our few-shot sampling outperform standard community models in inductive evaluation.

2 Related Work

2.1 Community Models

Community models have shown promise on static social graphs. Such models use social graphs to contextualise content by the users that interact with them, phrasing the detection task as node classification. Mishra et al. (2018, 2019) and Monti et al. (2019) find that GNNs over heterogeneous user–tweet graphs outperform models using only text or user-based features. Chandra et al. (2020) show that relational GNNs—which directly model edge relations between different types of nodes—yield significant improvements over a range of baselines.

Shu et al. (2019b) argue for the inclusion of publishers as another node type, with Ren et al. (2021) also including topics.

Nguyen et al. (2022) utilise temporal replies to model the dynamic user–content interactions. Temporal representations aid in early detection of malicious content (Jian Cui et al., 2021; Song et al., 2021a). However, they still assume static content. Others have focused on directly detecting actors posting the content (Tan et al., 2023; Mehta et al., 2022); we explicitly exclude actor modelling from our methodology since it mandates different ethical considerations (Mishra et al., 2021).

Gong et al. (2023) provide a review of graph representation learning for malicious content detection. They also conclude that cross-domain generalisation remains an understudied problem for graph-based malicious content detection.

2.2 Generalisable Content-only Models

Developing malicious content detectors for cross-domain generalisation is receiving increased attention. For example, many task-aware domain adaptation approaches have been proposed (Zhang et al., 2020, 2021; Mosallanezhad et al., 2022; Lin et al., 2022; Yue et al., 2022). These methods are either “aware” of the distribution of representations in different datasets, or use external models to correct representations post-hoc. Generating representations that are invariant to domain shifts is a related direction (Ding et al., 2022; Huang et al., 2023).

Utilising (large) language models on unseen texts has also shown promising results (Lee et al., 2021b; Chiu et al., 2022; AlKhamissi et al., 2022). Lee et al. (2021c) use multitask fine-tuning on RoBERTa (Liu et al., 2019), and show that few-shot adaptation on related, but unseen datasets improves performance over fine-tuning on individual tasks. Yue et al. (2023) specifically train content-only misinformation detectors to rapidly adapt. We, however, focus solely on community models for malicious content on social graphs, and to the best of our knowledge, are the first to do so.

2.3 Subgraph Sampling & Meta-learning

Closely related to the idea of rapid adaptation to new tasks and domains is the field of meta-learning. Herein, models are trained to optimise themselves using a minimal amount of examples. In NLP, this has been investigated for document (Yu et al., 2018; van der Heijden et al., 2021), sentence (Dou et al., 2019; Bansal et al., 2020), and token-level tasks

Dataset	GossipCop	CoAID	TwitterHateSpeech
Task	Rumour	Fake News	Hate Speech
Domain	Celebrity Gossip	COVID-19	Entertainment
Label Proportions	True Fake 77.12% 22.88%	True Fake 94.72% 5.28%	Racism Sexism None 11.97% 19.43% 68.60%
Doc-User Interaction	Retweet	Retweet	Authorship
# Documents	17 617	947	16 201
# Users	29 229	4 059	1 875
# Edges	2 334 554	61 254	65 600

Table 1: Statistics of graph datasets post-processing. Further details are in Appendix B.

(Holla et al., 2020). For a thorough review, we refer the reader to Lee et al. (2021a, 2022).

To operationalise meta-learning from subgraphs, Huang and Zitnik (2021) propose G-Meta. They assume local subgraphs preserve information of a larger graph, such that training a GNN on relevant subgraphs can induce rapid adaptation from limited context. Other graph meta-learning procedures exist, however, they do not utilise episodes of local subgraphs. We refer the reader to Mandal et al. (2021) and Zhang et al. (2022) for a comprehensive review of the field.

3 Datasets & Tasks

We use three widely-adopted social graph datasets to train and evaluate community models. Social graph datasets are difficult to collect and degrade as users or content are moderated out. The first dataset, GossipCop, is used for pre-training. The other two datasets, CoAID and TwitterHateSpeech, are reserved for evaluating generalisation to unseen graphs. Table 1 provides some statistics, which are further complemented by Appendix B. All datasets were rehydrated, i.e., rebuilt through the API, using the Twitter Academic API prior to May 2023. See ‘Redistribution of Twitter Content’.

GossipCop is one of two datasets introduced by Shu et al. (2019a). It comprises 20k fact-checked celebrity rumour articles, and around 500k interacting Twitter users. Labels correspond to the (now defunct) GossipCop fact-checking scores, covering a variety of (usually unreliable) publishers. Articles from a single trusted source, E!Online, were included to reduce class imbalance. Users are connected to articles and other users.

CoAID contains articles from the first months of the COVID-19 pandemic, collected by Cui and Lee

(2020). We omit the “social media” category as most contain short, poorly formatted text. Fake news articles are labelled using a variety of fact-checking websites, whereas truthful news comes from (unverified) mainstream media outlets. After rehydration, this dataset is substantially smaller than when originally devised, with most lost documents corresponding to the fake class. Users are connected to articles and other users.

TwitterHateSpeech differs in task, domain, and relational schema from the other datasets. Document nodes are tweets generated by Twitter users during a few seed events. Waseem and Hovy (2016) identified prolific hate speech tweeters, and include their followers and followees in the graph. They manually labelled all tweets as racist, sexist or innocuous (none). Especially racist (0.3%) users are over-represented, leading to few diverse regions in the graph. User-document connections indicate authorship (as opposed to tweet/retweet interactions). Users are also connected to other users.

4 Methodology

A social graph, $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, consists of a set of nodes \mathcal{V} and a set of edges \mathcal{E} indicating which nodes are incident to each other. A node’s r -radius neighbourhood $\mathcal{N}_r(v)$ contains all other nodes that can be reached by paths of length r (also called ‘hops’, a series of incident nodes), and always includes v .

All datasets considered require modelling two node-types: documents ($\mathcal{V}^{\text{docs.}}$) and users ($\mathcal{V}^{\text{(users)}}$). Hence, the social graph \mathcal{G} is heterogeneous. Document nodes v contain exogenous features x_v (i.e. the content representation), and have target labels $y_v \in \mathcal{Y}$. Documents are only connected to users ($\forall u \in \mathcal{N}_1(v), u \in \mathcal{V}^{\text{(users)}}$), whereas users may

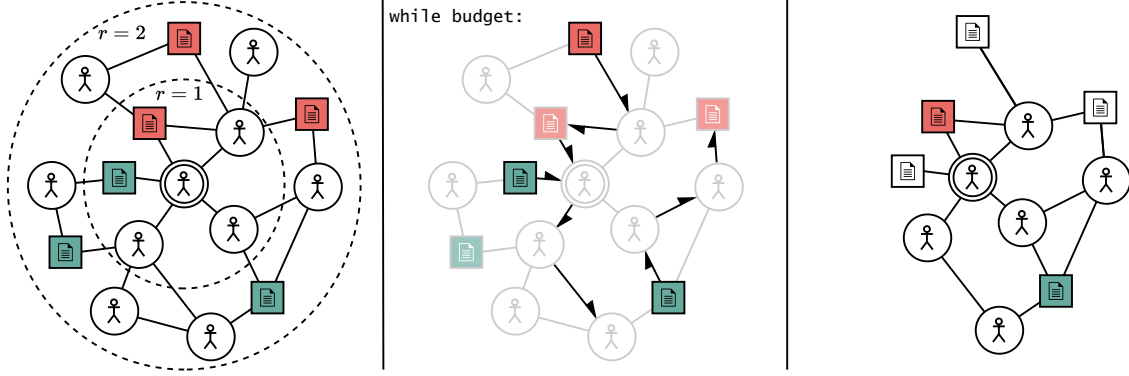


Figure 1: Support subgraph generation. Left: collect the r -radius neighbourhood of an anchor user. Middle: sub-sample using random walks from document nodes until reaching a maximum node count. Right: unmask document nodes inversely proportional to the number of subgraphs they appear in. Colours correspond to classes.

also be connected with other users based on their interactions or relations. Users are not labelled, and have no initial representation.

4.1 Community Modelling

Community models for malicious content detection classify content nodes in a social graph, taking into account the graph context around them to make the prediction $f_\theta(x_v; \mathcal{G})$. GNNs, a common representation learning framework, perform this contextualisation using non-linear message-passing schemes.

Some node v , at layer l , has as hidden state an aggregation of the representations in neighbouring nodes at layer $l - 1$. We use Graph Attention Networks (GATs) (Veličković et al., 2018), which employ an additive attention mechanism for the neighbourhood aggregation. Specifically:

$$\mathbf{h}_v^l = \sigma\left(\sum_{u \in \mathcal{N}_1(v)} \alpha_{v,u} \mathbf{W} \mathbf{h}_u^{l-1}\right) \quad (1)$$

where σ is a non-linear activation function. The attention weights $\alpha_{v,u}$ are computed as:

$$\text{softmax}\left(\sigma\left(\mathbf{a}^\top [\mathbf{W} \mathbf{h}_v^{l-1} \parallel \mathbf{W} \mathbf{h}_u^{l-1}]\right)\right) \quad (2)$$

where $[\cdot \parallel \cdot]$ is concatenation, while \mathbf{a} and \mathbf{W} are the projection matrices.

More expressive architectures than GATs exist and have been applied to malicious content detection (Chandra et al., 2020; Gong et al., 2023). Such models, however, often introduce inductive biases specific to the task it seeks to solve. For example, relational attention aids performance but requires the relational schema to be consistent across datasets. Our evaluation and meta-learning setup is model agnostic.

Community modelling can be *transductive* or *inductive*. Transductive modelling assumes that

the social graph remains static across training and prediction. Inductive modelling, instead, assumes the underlying social graph changes, in terms of content and users.

Herein we differ from the definition common to graph learning applications. Usually, inductive graph learning ‘only’ assumes the nodes in the evaluation graph are unseen, with those nodes coming from the same underlying graph. As argued in the introduction, this does not apply to malicious content detection; the graph has shifted between training and deployment time. True inductive generalisation, therefore, requires generalisation to entirely different graphs. Currently, no graph datasets of malicious content exist that allow testing this manner of generalisation.

Due to the social network changing, inductive community models should not rely on superficial properties, like the content of malicious posts or specific user neighbourhoods, but rather leverage universal network properties. One such property is homophily: the tendency of nodes of a similar class to cluster together. We investigate the presence of homophily (or heterophily) in Appendix F.

4.2 Few-shot Subgraph Sampling

A successful community model for malicious content detection should be able to rapidly adapt to the constantly evolving social graph, even when presented with labelled examples.

More formally, a community model, f_θ , should be able to inductively learn to generalise from a limited exploration of a social graph $\mathcal{G}^S \subset \mathcal{G}^I$ to make accurate predictions elsewhere $\mathcal{G}^Q = \mathcal{G}^I - \mathcal{G}^S$. In commonly-used meta-learning terminology, S would denote the support and Q the query set.

Algorithm 1 Few-shot Graph Sampling

Require: Graph $\mathcal{G} = ((\mathcal{V}^{(\text{doc})}, \mathcal{V}^{(\text{user})}), \mathcal{E})$
Require: Central user v
Require: Maximum support graph size, budget

for r in range(2, 5) **do**
 Get r -radius neighbourhood of v , $\mathcal{N}_r(v)$
 Find nodes per class $\mathcal{V}_y^{(\text{doc})}$
 if all($|\mathcal{V}_y^{(\text{doc})}| \geq k$ -shot) **then**
 break
 end if
end for

Initialize empty support graph \mathcal{G}'
while $|\mathcal{G}'| < \text{budget}$ **do**
 Pick class y round-robin style
 Sample a node u from $\mathcal{V}_y^{(\text{doc})}$
 Generate random walk path, $\mathcal{G}^{(\text{path})}$, from u
 $\mathcal{G}' \leftarrow \mathcal{G}' \cup \mathcal{G}^{(\text{path})}$
end while

return \mathcal{G}' if it is a valid k -shot graph

For malicious content detection specifically, the notion of ‘limited exploration’ implies the following conditions for \mathcal{G}^S :

1. **Locality:** all sampled document nodes come from the same graph region, due to a similar seed event, topic, or intended audience
2. **Limited Context:** moderation should precede wide-spread dissemination
3. **Few-shot:** labelling is expensive, therefore a minimal set of labels is available

Existing subgraph sampling procedures, like G-Meta, violate these conditions, especially ‘locality’. Labelled nodes are sampled independently, i.e., nodes can come from entirely different regions of the graph, which in our case, would imply entirely unrelated forms of content.

To better conform to the listed desiderata, we perform user-centred sampling for generating \mathcal{G}^S . Algorithm 1 presents pseudocode for our sampling approach, which is also depicted graphically in Figure 1. Various graph statistics are provided in Appendix C.

To ensure locality, we first sample an anchor user and collect the smallest r -hop neighbourhood that yields k documents of each class. In Figure 1, the double-circled user represents the anchor. Then, to

limit social context, we take random walks from the documents nodes into the subgraph. This process starts from a document node, and is repeated until a maximum number of nodes is reached. Bold arrows in the middle column of Figure 1 show some random walks of length 3.

For the training process, only k document nodes of each class have their labels unmasked in a subgraph. Other document nodes are still allowed in the subgraph, but without labels. This is depicted in the right-most column of Figure 1.

Power-law distributed degrees of nodes is a property of social media networks. This means a few, very active users and their incident document nodes will be present in the majority of subgraphs. This reduces the diversity of support episodes and thus biases generalisability estimates. To reduce the effect of these users during the training process, document nodes are labelled inversely proportional to their frequency across all created subgraphs.

4.3 Gradient-based Meta-learning

Community models learn a neighbourhood-aware mapping of a content node’s input features to target labels. Community meta-learners, instead, use an initial set of weights to produce community models only after adaptation, i.e. learning a community model from several episodes of \mathcal{G}^S and \mathcal{G}^Q . By using our few-shot subgraph sampling method to create episodes for meta-training, the community meta-learners are better suited to inductive generalisation.

We focus on a specific subclass of meta-learners, namely, gradient-based meta-learners. Model-Agnostic Meta-Learning (MAML), introduced by Finn et al. (2017), is the most popular such learning framework. Its optimisation objective is:

$$\min_{\theta^{(\text{meta})}} \mathbb{E}[\mathcal{L}(\mathbf{y}_Q, f_{\theta^{(\text{task})}}(\mathbf{x}_Q; \mathcal{G}^Q))] \quad (3)$$

The induced update to $\theta^{(\text{meta})}$ is called the *outer-loop* update. The *inner-loop* occurs during adaptation to the support set, using a pre-defined number of SGD updates, $t \in \{1, \dots, T_{\text{inner}}\}$, with gradients

$$\nabla_{\theta_t^{(\text{task})}} \mathcal{L}(\mathbf{y}_S, f_{\theta_t^{(\text{task})}}(\mathbf{x}_S; \mathcal{G}^S)) \quad (4)$$

This bi-level objective encourages the meta-learner’s initial weights, $\theta^{(\text{meta})}$, to learn to adapt to new tasks, $\theta^{(\text{task})}$, using only T_{inner} updates.

Prototypical Initialisation MAML implicitly assumes a new permutation of classes in each episode and re-initialises the task-specific classification head during each outer-loop iteration. Prototypical Networks (ProtoNets) (Snell et al., 2017) are a non-gradient-based meta-learning alternative that does not utilise classification heads. Instead, support samples are used to form class prototypes \mathbf{c}_y :

$$\mathbf{c}_y = \frac{1}{k} \sum_{\{\mathbf{x}_v | y_v = y\}} f_{\theta(\text{meta})}(\mathbf{x}_S; \mathcal{G}^S) \quad (5)$$

which classify query samples based on their distance to the prototypes.

$$p(y_v | \mathbf{x}_v) = \text{softmax}(-d(f_{\theta(\text{meta})}(\mathbf{x}_v), \mathbf{c}_y)) \quad (6)$$

Per Triantafillou et al. (2020), if using the Euclidean distance as d , this is equivalent to applying a linear projection $\mathbf{W}\mathbf{h} + \mathbf{b}$ with initialisation:

$$\mathbf{W} = 2\mathbf{c}_y, \quad \mathbf{b} = \|\mathbf{c}_y\|^2 \quad (7)$$

Using this reformulation, Triantafillou et al. (2020) propose ProtoMAML, an approach that extends MAML such that the classification head is now parameterised by Eq. 7 and fully updatable during inner loop adaptation.

4.4 Implementation Details

Support graphs for episodes are sampled using the few-shot sampling procedure detailed in Sec. 4.2. We use the lowest possible radius r that satisfies the k -shot requirement. The maximum number of nodes in the support graph is 2048. All classes provide an equal amount of root nodes of the random walk sub-sampling. For meta-training, $k = 4$, and the query graph is generated by sampling a set of independent document nodes along with their $r = 2$ neighbourhood. During (meta-)testing, all non-training nodes are used.

Document nodes use the time-pooled averaged token representations from the penultimate RoBERTa (Liu et al., 2019) layer as initial representations. These are not trained further. Users nodes are initialised to all zeros, making them effectively anonymous and allowing for both transductive and inductive approaches. All models train end-to-end and do not include an auxiliary text-only classifier. Appendix D provides all additional modelling hyper-parameters.

Our GNN architecture is adapted from SAFER (Chandra et al., 2020). It consists of 2 ReLU activated GAT layers, each with 3 attention heads.

These are concatenated together and linearly projected before being fed into a 2-layer MLP. We use dropout node-wise on the initial representations and element-wise on the layer representation and attention weights. We reduce the computational complexity of the GAT layers by merging successive projections in the attention layers (Brody et al., 2022). The use of 2 GAT layers means document nodes v have as receptive field $\mathcal{N}_2(v)$. We optimise our models using AdamW (Kingma and Ba, 2017; Loshchilov and Hutter, 2019).

In total, we experiment with models trained under 6 different learning paradigms. The first two (FULL and SUBGRAPHS) are non-episodic baselines, trained transductively on the full graph or inductively on few-shot sampled subgraphs respectively. FULL mimics the current practice of training transductively without generalisation to unseen graphs in mind. SUBGRAPHS makes generalisation feasible and allows us to isolate the contribution of meta-learning.

The last four are graph meta-learners. We use two MAML variants, one with a classification head shared across episodes (MAML-LH) and another where the classification head is randomly initialised at each episode (MAML-RH). Appendix E shows the effect of classifier head resetting on adaptation speed. We also train PROTONET and PROTOMAML variants to evaluate the effect of prototypical initialisation on the classification head.

MAML-based outer-loop updates (Eq. 3) require computing second-order gradients, which is prohibitively expensive. Instead, we use a first-order approximation (foMAML (Finn et al., 2017)) which usually does not significantly affect performance (Nichol et al., 2018; Antoniou et al., 2019).

5 Experiments and Results

5.1 Experimental Setup

We first assess within-dataset generalisation to unseen nodes using 5-fold stratified cross-validation. The folds are strict, with no document appearing in more than one validation set. We only keep the largest connected component. User nodes can appear in multiple folds, but since they are all zero-initialised, they cannot influence nodes in other folds.

Models are trained on GossipCop and then evaluated on all three datasets. On GossipCop, we assess both the non-episodic and episodic models. On other datasets, we assess the episodic models only

Method		F1		AUPR	MCC
		Real	Fake		
SAFER (Chandra et al., 2020)			0.9453		
Baselines	TEXT	0.8773 (0.8628, 0.8918)	0.5963 (0.5854, 0.6072)	0.6738 (0.6570, 0.6905)	0.4767 (0.4532, 0.5002)
	USER ID	0.9423 (0.9403, 0.9444)	0.7431 (0.729, 0.7572)	0.8644 (0.8556, 0.8733)	0.7164 (0.7051, 0.7277)
	FULL	<u>0.9672</u> (0.9615, 0.9728)	<u>0.8920</u> (0.8754, 0.9086)	0.9450 (0.9291, 0.9608)	<u>0.8601</u> (0.8384, 0.8818)
	SUBGRAPHS	0.9485 (0.9406, 0.9563)	0.8496 (0.8326, 0.8666)	<u>0.9473</u> (0.9412, 0.9534)	0.8088 (0.7886, 0.8289)
GAT	MAML-LH [†]	0.9732 (0.9731, 0.9732)	0.9092 (0.9091, 0.9094)	0.9651 (0.9651, 0.9651)	0.8828 (0.8826, 0.8831)
	MAML-RH [†]	0.8861 (0.8776, 0.8946)	0.7559 (0.7498, 0.7620)	0.9164 (0.9136, 0.9192)	0.7108 (0.7021, 0.7194)
	PROTONET [†]	0.9205 (0.9121, 0.9289)	0.8192 (0.8116, 0.8268)	0.9175 (0.9099, 0.9250)	0.7535 (0.7384, 0.7686)
	PROTOMAML [†]	0.8921 (0.8825, 0.9018)	0.7925 (0.7857, 0.7994)	0.9255 (0.922, 0.9290)	0.7263 (0.7158, 0.7369)

Table 2: Results on GossipCop. Brackets give the 90% confidence interval. Bold values denote the best column score (where comparison is possible) and underlined the second-best. [†] denotes 4-shot episodic evaluation. SAFER results taken from Chandra et al. (2020); FULL is our re-implementation since they do not release their data splits.

to ensure a few-shot generalisation setup. Within each episode, support nodes appear in the query graph but do not count towards classification performance metrics. This process is repeated 256 times, with summary statistics computed for each of the 5-fold model checkpoints. When aggregating over checkpoints, the inverse-variance weighted mean is used to estimate a common effect size (i.e., a fixed-effect meta-analysis (Schwarzer et al., 2015)).

To assess classification performance for each class in isolation, we use the F1-score. Matthews Correlation Coefficient (MCC) is used to assess holistically. Recent papers argue for the MCC as an informative metric, relatively robust to class imbalance (Chicco and Jurman, 2020, 2022). MCC values near 0 indicate random performance, values near 1 almost perfect performance, and negative values are worse than random. The Area Under the Precision-Recall curve (AUPR) is a multi-threshold metric, and can compare models on their ability to separate classes. We exclude it for CoAID and TwitterHateSpeech as there is no consistent way to do aggregate it in multi-class settings. Metrics are reported with 90% confidence intervals.

All hyper-parameters used were tuned on GossipCop’s validation sets. The tuning procedure, optimizer, meta-learning and evaluation hyper-parameters are described in Appendix D.

5.2 GossipCop Results

Here, we test generalisation on unseen nodes from the same graph. Beyond the non-episodic baselines already described, we have two additional baseline methods on GossipCop. The first, TEXT, is a 2 layer MLP on top of the initial document embeddings meant to test the added benefit of a graph inductive bias. The second, USER ID, classifies test documents based on neighbouring users’ most linked document class in the train split. The already reasonable performance indicates high homophily.

The GAT-based models leverage both text and social features. SUBGRAPHS clearly performs worse than FULL. MAML-LH, however, outperforms FULL even though it is inductive, demonstrating the generalisation power of meta-learning.

The lower three rows all include meta-learners which constantly re-initialise the classification head. Their performance is more in line with the non-episodic SUBGRAPHS, lagging considerably on fixed threshold metrics. This gap narrows in terms of AUPR, implying the final bias parameter may be to blame.

5.3 Generalisation to Unseen Graphs

Figure 2 shows the performance of models ported to the other two datasets.

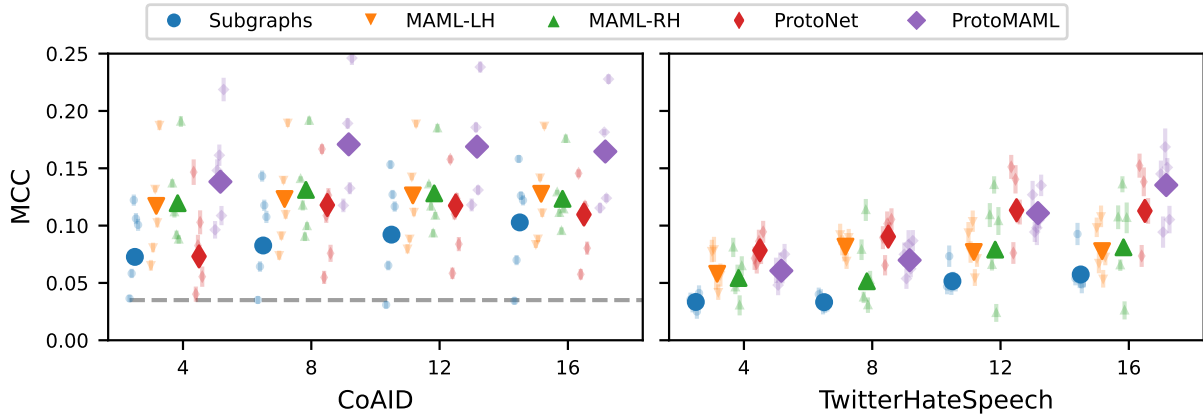


Figure 2: Generalisation of various models to CoAID and TwitterHateSpeech, in terms of MCC. See-through markers give the performance of each model instance, with error bars giving the 90% CI. Solid markers give the performance averaged across model instances. Markers are offset to avoid overlap. The horizontal axis gives the support graph k -shot. The dashed gray line for CoAID gives the zero-shot performance of the Subgraphs model, i.e., direct domain transfer. Colours and shape both denote a model instance.

Variance between the different model instances is large, although when aggregated their performance is relatively stable. PROTOMAML proves to be the best model on both datasets, particularly at larger k -shot values. PROTONET shows competitive performance on TwitterHateSpeech (especially in lower k -shot settings), but is considerably worse on CoAID. Prototypical initialisation seems to aid generalisation, mitigating classification head learning. Regardless, meta-learning methods outperform the non-episodic SUBGRAPHS model on either dataset, indicating that training for rapid adaptation helps generalisation to new malicious content forms.

Transfer to CoAID from Gossipcop, is essentially a form of domain transfer. As such, we provide the zero-shot performance of the SUBGRAPHS model as a baseline value. Despite the similar task definition, adaptation is clearly required for generalisation, as evidenced by the near-random performance of SUBGRAPHS in the zero-shot setting, and the aggressive hyperparameters required (see Appendix G.1).

Table 3 provides F1 scores for PROTOMAML on each class. All-in-all, the highest achieved MCC was 0.1709, for PROTOMAML at $k = 8$, corresponding to an F1-Fake of 0.1841. While low relative to other F1 scores reported, this should be compared to a class prevalence of 5%.

Relative to random performance, the greatest negative outlier is TwitterHateSpeech’s majority class, ‘None’. One possible explanation is the homophily pattern of TwitterHateSpeech (see Appendix F). Whereas racist and sexist tweets are

primarily homophilic in the query set, a large proportion of innocuous tweets are highly *heterophilic*; i.e. these are contextualised by users predominantly authoring racist or sexist. The model is therefore more likely to err on those innocuous tweets, as their author shows a proclivity towards hate speech. Here, heterophily serves as noise. This is most likely an artefact of Waseem and Hovy’s data collection process, with prolific racist and sexist authors serving as the anchor around which the rest of the graph is built.

In general, the results here do not correlate with those found in Table 2. Underperformers there show relatively better performance after adaptation to the other datasets. Hinting at overfitting to the GossipCop graph, this aligns with the line of argumentation presented in the introduction: performance on a single, static graph is not indicative of generalisation to emerging malicious content.

5.4 Ablating GossipCop Pre-Training

To test the effect of GossipCop pre-training on generalisation to other datasets, we ablate PROTOMAML’s pre-trained weights, and repeat the evaluation under re-initialised weights. A comparison in terms of MCC is provided in Table 4. On CoAID, PROTOMAML outperforms PROTOMAML-RESET at all k -shot values. On TwitterHateSpeech, this only happens at the larger k -shot values, with PROTOMAML’s MCC performance increase outpacing its RESET counterpart.

While low, comparing the performance on each class individually (Tables 3 and 15), PROTOMAML

k	CoAID		TwitterHateSpeech		
	Real	Fake	Racist	Sexist	None
4	0.7734	0.1762	0.1763	0.2181	0.3585
8	0.8245	0.1841	0.1934	0.2148	0.3530
12	0.8245	0.1732	0.2545	0.2554	0.3503
16	0.8321	0.1599	0.3021	0.3077	0.3163
B	0.6545	0.0955	0.1932	0.2799	0.5784

Table 3: F1 scores achieved by PROTOMAML during generalisation to the auxiliary datasets. Row B provides F1 scores for a random classifier (Flach and Kull, 2015). This table is complemented by Appendix G.

is able to increase its performance on all classes simultaneously, whereas PROTOMAML-RESET only does so for the racist class, degrading performance on sexist and innocuous documents.

Regardless, the fact that a model trained specifically with generalisation in mind is barely able to outperform one with random weights is striking, and speaks to the inadequacy of GossipCop as an evaluation dataset. Furthermore, near perfect performance on unseen nodes of the pre-training graph does not imply inductive generalisation to new graphs.

6 Conclusion

This paper proposes a more realistic evaluation setup for community models on malicious content detection. We highlight several properties of evolving social graphs that are especially neglected: expensive labelling, limited context, and emerging content and users. Experiments verified our motivation, with performance on a single, static dataset in a transductive manner bearing little resemblance to performance during few-shot inductive generalisation.

Our proposed few-shot subgraph sampling approach presented in Section 4.2 is tailored to social media graphs and allows generalisation of community models to new networks, domains, and tasks. While standard community models performed poorly, incorporating our sampling procedure in graph meta-learners aided generalisation. Particularly promising are models with prototypical initialisation.

Ultimately, our results suggest that malicious content detection using community models is not ‘solved’, despite some models achieving near perfect evaluation scores. Current evaluation proce-

k	CoAID		TwitterHateSpeech	
	Trained	Reset	Trained	Reset
4	0.1383	0.1191	0.0607	0.0767
8	0.1709	0.1398	0.0699	0.0868
12	0.1689	0.1304	0.1109	0.1025
16	0.1646	0.1212	0.1354	0.1052
B	0.0000	0.0000	0.0000	0.0000

Table 4: MCC scores achieved by PROTOMAML (Trained) and PROTOMAML-RESET (Reset) during generalisation to the auxiliary datasets. This table is complemented by Appendix G.3.

dures neglect critical properties of malicious content, and models tested under these conditions will not prove useful in realistic deployment settings. This is a regrettable consequence, considering the high-stakes nature of malicious content detection. Much like the trend occurring in the content-only malicious content detection literature (see Section 2.2), we hope this work will lead to similar follow-up work for community models.

An open problem, warranting further investigation, is the application of meta-learning to class imbalanced datasets. Common to malicious content detection, class imbalance imposes a severe penalty on meta-learners that reset their classification heads.

7 Limitations

While we took care to increase the diversity of the training data (user-centred sampling, distributing the labels, adding dropout throughout models), ultimately, the diversity is limited by the underlying graph dataset. GossipCop is large, but it contains only a single task and a relatively uniform structure. Ideally, multiple, distinct graph datasets would be used in meta-training. However, few such datasets are available. For meta-learning, task diversity might be a critical factor in ‘learning-to-learn’, analogous to data diversity being critical in standard machine-learning setups. As such, our meta-learners are likely operating below capacity.

Phan et al. (2023) made the conclusion that no common benchmarks for community models of malicious content are currently in use. This work, despite making a step towards more realistic evaluation of such models, does not improve this situation. The presented performance metrics are in line with related work, but meaningful comparison will

only be possible with the publication and adoption of open-access benchmark datasets. This seems an unrealistic short-term aspiration at the time of writing.

8 Ethical Considerations

Following the work of [Mishra et al. \(2021\)](#), we take some steps to ensure that our experimental setup addresses the ethical considerations that may arise when modelling users and communities. The authors highlight the following three considerations that apply to our work:

- *Personal vs. Population-level trends*, i.e., are generalisations being made from personal traits to population-level trends
- *Bias in datasets*, i.e., is there demographic, comment distribution, or label bias in the dataset(s) being used?
- *Purpose*, i.e., is the purpose of the modelling being done to classify content as malicious or users and communities too?

In order to tackle the *comment distribution* bias whereby the majority of documents may belong to a small number of users, we remove users with more than a certain number of documents from the dataset (where possible). Furthermore, to counter the *label distribution* bias where we only pick documents of a particular class from a specific user, we do user-centred sampling, incorporating the entire neighbourhood of a user in the user-document graph. We initialise all users to the same zero-embedding, ensuring that we do not generalise personal traits to population-level trends. Lastly, in our work, we solely leverage the user-document graph to be able to better classify the documents, not the users themselves as malicious, hence having a clear purpose to advance malicious content detection.

References

- Lada A. Adamic, Thomas M. Lento, Eytan Adar, and Pauline C. Ng. 2016. [Information Evolution in Social Networks](#). In *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining*, pages 473–482.
- Raed Alharbi, Minh N. Vu, and My T. Thai. 2021. [Evaluating Fake News Detection Models from Explainable Machine Learning Perspectives](#). In *ICC 2021 - IEEE International Conference on Communications*, pages 1–6.
- Badr AlKhamissi, Millicent Li, Asli Celikyilmaz, Mona Diab, and Marjan Ghazvininejad. 2022. [A Review on Language Models as Knowledge Bases](#).
- Hunt Allcott and Matthew Gentzkow. 2017. [Social media and fake news in the 2016 election](#). *Journal of Economic Perspectives*, 31(2):211–36.
- Antreas Antoniou, Harrison Edwards, and Amos Storkey. 2019. [How to train your MAML](#).
- Trapit Bansal, Rishikesh Jha, and Andrew McCallum. 2020. [Learning to Few-Shot Learn Across Diverse Natural Language Classification Tasks](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5108–5123, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Lia Bozarth and Ceren Budak. 2020. [Toward a Better Performance Evaluation Framework for Fake News Classification](#). *Proceedings of the International AAAI Conference on Web and Social Media*, 14:60–71.
- Shaked Brody, Uri Alon, and Eran Yahav. 2022. [How Attentive are Graph Attention Networks?](#)
- Shantanu Chandra, Pushkar Mishra, Helen Yanakoudakis, Madhav Nimishakavi, Marzieh Saeidi, and Ekaterina Shutova. 2020. [Graph-based Modeling of Online Communities for Fake News Detection](#).
- Davide Chicco and Giuseppe Jurman. 2020. [The advantages of the Matthews correlation coefficient \(MCC\) over F1 score and accuracy in binary classification evaluation](#). *BMC Genomics*, 21(1):6.
- Davide Chicco and Giuseppe Jurman. 2022. [An Invitation to Greater Use of Matthews Correlation Coefficient in Robotics and Artificial Intelligence](#). *Frontiers in Robotics and AI*, 9.
- Ke-Li Chiu, Annie Collins, and Rohan Alexander. 2022. [Detecting Hate Speech with GPT-3](#).
- Limeng Cui and Dongwon Lee. 2020. [CoAID: COVID-19 Healthcare Misinformation Dataset](#).
- Carner Derron. 2021. [Commission on Information Disorder Final Report](#). <https://www.aspeninstitute.org/publications/commission-on-information-disorder-final-report/>.
- Yasan Ding, Bin Guo, Yan Liu, Yunji Liang, Haocheng Shen, and Zhiwen Yu. 2022. [MetaDetector: Meta Event Knowledge Transfer for Fake News Detection](#). *ACM Transactions on Intelligent Systems and Technology*.
- Zi-Yi Dou, Keyi Yu, and Antonios Anastasopoulos. 2019. [Investigating Meta-Learning Algorithms for Low-Resource Natural Language Understanding Tasks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages

- 1192–1197, Hong Kong, China. Association for Computational Linguistics.
- Brussels European Commission. 2018. [Flash eurobarometer 464 \(fake news and disinformation online\)](#). GESIS Data Archive, Cologne. ZA6934 Data file Version 1.0.0, <https://doi.org/10.4232/1.13019>.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. [Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks](#). *CoRR*.
- Peter Flach and Meelis Kull. 2015. Precision-Recall-Gain Curves: PR Analysis Done Right. In *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc.
- Shuzhi Gong, Richard O. Sinnott, Jianzhong Qi, and Cecile Paris. 2023. [Fake News Detection Through Graph-based Neural Networks: A Survey](#).
- Mingfei Guo, Xiuying Chen, Juntao Li, Dongyan Zhao, and Rui Yan. 2021. [How does Truth Evolve into Fake News? An Empirical Study of Fake News Evolution](#). In *Companion Proceedings of the Web Conference 2021*, pages 407–411, Ljubljana Slovenia. ACM.
- Nithin Holla, Pushkar Mishra, Helen Yannakoudakis, and Ekaterina Shutova. 2020. [Learning to Learn to Disambiguate: Meta-Learning for Few-Shot Word Sense Disambiguation](#).
- Benjamin Horne. 2019. [NELA2017](#).
- Benjamin Horne and Mauricio Gruppi. 2021. [NELA-GT-2020](#).
- Benjamin Horne and Mauricio Gruppi. 2023. [NELA-GT-2022](#).
- Kexin Huang and Marinka Zitnik. 2021. [Graph Meta Learning via Local Subgraphs](#).
- Yinqiu Huang, Min Gao, Jia Wang, Junwei Yin, Kai Shu, Qilin Fan, and Junhao Wen. 2023. [Meta-prompt based learning for low-resource false information detection](#). *Information Processing & Management*, 60(3):103279.
- Hussain Hussain, Tomislav Duricic, Elisabeth Lex, Denis Helic, and Roman Kern. 2021. [The interplay between communities and homophily in semi-supervised classification using graph neural networks](#). *Applied Network Science*, 6(1):1–26.
- Jian Cui, Kwanwoo Kim, Seung Ho Na, and Seungwon Shin. 2021. [Hetero-SCAN: Towards Social Context Aware Fake News Detection via Heterogeneous Graph Neural Network](#). *arXiv: Social and Information Networks*.
- Diederik P. Kingma and Jimmy Ba. 2017. [Adam: A Method for Stochastic Optimization](#).
- Hung-yi Lee, Shang-Wen Li, and Ngoc Thang Vu. 2022. [Meta Learning for Natural Language Processing: A Survey](#).
- Hung-yi Lee, Ngoc Thang Vu, and Shang-Wen Li. 2021a. [Meta Learning and Its Applications to Natural Language Processing](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: Tutorial Abstracts*, pages 15–20, Online. Association for Computational Linguistics.
- Nayeon Lee, Yejin Bang, Andrea Madotto, and Pascale Fung. 2021b. [Towards Few-shot Fact-Checking via Perplexity](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1971–1981, Online. Association for Computational Linguistics.
- Nayeon Lee, Belinda Z. Li, Sinong Wang, Pascale Fung, Hao Ma, Wen-tau Yih, and Madian Khabsa. 2021c. [On Unifying Misinformation Detection](#).
- Derek Lim, Felix Hohne, Xiuyu Li, Sijia Linda Huang, Vaishnavi Gupta, Omkar Bhalerao, and Ser-Nam Lim. 2021. [Large Scale Learning on Non-Homophilous Graphs: New Benchmarks and Strong Simple Methods](#).
- Hongzhan Lin, Jing Ma, Liangliang Chen, Zhiwei Yang, Mingfei Cheng, and Chen Guang. 2022. [Detect Rumors in Microblog Posts for Low-Resource Domains via Adversarial Contrastive Learning](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 2543–2556, Seattle, United States. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [RoBERTa: A Robustly Optimized BERT Pretraining Approach](#).
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled Weight Decay Regularization](#).
- Yao Ma, Xiaorui Liu, Neil Shah, and Jiliang Tang. 2021. [Is Homophily a Necessity for Graph Neural Networks?](#)
- Debmalya Mandal, Sourav Medya, Brian Uzzi, and Charu Aggarwal. 2021. [MetaLearning with Graph Neural Networks: Methods and Applications](#). *ACM SIGKDD Explorations Newsletter*, 23(2):13–22.
- Nikhil Mehta, Maria Leonor Pacheco, and Dan Goldwasser. 2022. [Tackling fake news detection by continually improving social context representations using graph neural networks](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1363–1380, Dublin, Ireland. Association for Computational Linguistics.
- Pushkar Mishra, Marco Del Tredici, Helen Yannakoudakis, and Ekaterina Shutova. 2018. [Author Profiling for Abuse Detection](#). In *Proceedings of the*

- 27th International Conference on Computational Linguistics, pages 1088–1098, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Pushkar Mishra, Marco Del Tredici, Helen Yannakoudakis, and Ekaterina Shutova. 2019. [Abusive Language Detection with Graph Convolutional Networks](#).
- Pushkar Mishra, Aleksandra Piktus, Gerard Goossen, and Fabrizio Silvestri. 2020. [Node masking: Making graph neural networks generalize and scale better](#). *CoRR*, abs/2001.07524.
- Pushkar Mishra, Helen Yannakoudakis, and Ekaterina Shutova. 2021. [Modeling users and online communities for abuse detection: A position on ethics and explainability](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3374–3385, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Federico Monti, Fabrizio Frasca, Davide Eynard, Damon Mannon, and Michael M. Bronstein. 2019. [Fake News Detection on Social Media using Geometric Deep Learning](#).
- Ahmadreza Mosallanezhad, Mansooreh Karami, Kai Shu, Michelle V. Mancenido, and Huan Liu. 2022. [Domain Adaptive Fake News Detection via Reinforcement Learning](#). In *Proceedings of the ACM Web Conference 2022*, pages 3632–3640.
- Karsten Müller and Carlo Schwarz. 2017. [Fanning the flames of hate: Social media and hate crime](#). *SSRN Electronic Journal*.
- M. E. J. Newman. 2003. [Mixing patterns in networks](#). *Physical Review E*, 67(2):026126.
- Van-Hoang Nguyen, Kazunari Sugiyama, Preslav Nakov, and Min-Yen Kan. 2022. [FANG: Leveraging social context for fake news detection using graph representation](#). *Communications of the ACM*, 65(4):124–132.
- Alex Nichol, Joshua Achiam, and John Schulman. 2018. [On First-Order Meta-Learning Algorithms](#).
- Dan Saattrup Nielsen and Ryan McConville. 2022. [MuMiN: A Large-Scale Multilingual Multimodal Fact-Checked Misinformation Social Network Dataset](#).
- Jaehoon Oh, Hyungjun Yoo, ChangHwan Kim, and Se-Young Yun. 2021. [BOIL: Towards Representation Change for Few-shot Learning](#).
- Huyen Trang Phan, Ngoc Thanh Nguyen, and Dosam Hwang. 2023. [Fake news detection: A survey of graph neural network methods](#). *Applied Soft Computing*, 139:110235.
- Aniruddh Raghu, Maithra Raghu, Samy Bengio, and Oriol Vinyals. 2020. [Rapid Learning or Feature Reuse? Towards Understanding the Effectiveness of MAML](#).
- Yuxiang Ren, Bo Wang, Jiawei Zhang, and Yi Chang. 2021. [Adversarial Active Learning based Heterogeneous Graph Neural Network for Fake News Detection](#).
- Giancarlo Ruffo, Alfonso Semeraro, Anastasia Giachanou, and Paolo Rosso. 2023. [Studying Fake News Spreading, Polarisation Dynamics, and Manipulation by Bots: A Tale of Networks and Language](#). *Computer Science Review*, 47:100531.
- Guido Schwarzer, James R. Carpenter, and Gerta Rücker. 2015. [Meta-Analysis with R. Use R!](#) Springer International Publishing, Cham.
- Kai Shu, Deepak Mahudeswaran, Suhang Wang, Dongwon Lee, and Huan Liu. 2019a. [FakeNewsNet: A Data Repository with News Content, Social Context and Spatiotemporal Information for Studying Fake News on Social Media](#).
- Kai Shu, Suhang Wang, and Huan Liu. 2019b. [Beyond News Contents: The Role of Social Context for Fake News Detection](#). In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, pages 312–320, Melbourne VIC Australia. ACM.
- Jake Snell, Kevin Swersky, and Richard S. Zemel. 2017. [Prototypical Networks for Few-shot Learning](#).
- Chenguang Song, Kai Shu, and Bin Wu. 2021a. [Temporally evolving graph neural network for fake news detection](#). *Information Processing and Management: an International Journal*, 58(6).
- Zixing Song, Xiangli Yang, Zenglin Xu, and Irwin King. 2021b. [Graph-based semi-supervised learning: A comprehensive review](#). *CoRR*, abs/2102.13303.
- Ruoyu Sun, Cong Li, Barbara Millet, Khudejah Iqbal Ali, and John Petit. 2022. [Sharing news with online friends: A study of network homophily, network size, and news type](#). *Telematics and Informatics*, 67:101763.
- Zhaoxuan Tan, Shangbin Feng, Melanie Sclar, Herun Wan, Minnan Luo, Yejin Choi, and Yulia Tsvetkov. 2023. [Botpercent: Estimating bot populations in twitter communities](#).
- Eleni Triantafillou, Tyler Zhu, Vincent Dumoulin, Pascal Lamblin, Utku Evci, Kelvin Xu, Ross Goroshin, Carles Gelada, Kevin Swersky, Pierre-Antoine Manzagol, and Hugo Larochelle. 2020. [Meta-Dataset: A Dataset of Datasets for Learning to Learn from Few Examples](#).
- Niels van der Heijden, Helen Yannakoudakis, Pushkar Mishra, and Ekaterina Shutova. 2021. [Multilingual and cross-lingual document classification: A meta-learning approach](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1966–1976, Online. Association for Computational Linguistics.

- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. [Graph Attention Networks](#).
- Soroush Vosoughi, Deb Roy, and Sinan Aral. 2018. [The spread of true and false news online](#). *Science*, 359(6380):1146–1151.
- Zeeraq Waseem and Dirk Hovy. 2016. [Hateful Symbols or Hateful People? Predictive Features for Hate Speech Detection on Twitter](#). In *Proceedings of the NAACL Student Research Workshop*, pages 88–93, San Diego, California. Association for Computational Linguistics.
- Daniel Wiessner. 2021. Judge OKs \$85 mln settlement of Facebook moderators’ PTSD claims. *Reuters*.
- Mo Yu, Xiaoxiao Guo, Jinfeng Yi, Shiyu Chang, Saloni Potdar, Yu Cheng, Gerald Tesauro, Haoyu Wang, and Bowen Zhou. 2018. [Diverse Few-Shot Text Classification with Multiple Metrics](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1206–1215, New Orleans, Louisiana. Association for Computational Linguistics.
- Zhenrui Yue, Huimin Zeng, Ziyi Kou, Lanyu Shang, and Dong Wang. 2022. [Contrastive Domain Adaptation for Early Misinformation Detection: A Case Study on COVID-19](#). In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, pages 2423–2433, Atlanta GA USA. ACM.
- Zhenrui Yue, Huimin Zeng, Yang Zhang, Lanyu Shang, and Dong Wang. 2023. [MetaAdapt: Domain Adaptive Few-Shot Misinformation Detection via Meta Learning](#).
- Hanqing Zeng, Hongkuan Zhou, Ajitesh Srivastava, Rajgopal Kannan, and Viktor Prasanna. 2020. [GraphSAINT: Graph Sampling Based Inductive Learning Method](#).
- Chuxu Zhang, Kaize Ding, Jundong Li, Xiangliang Zhang, Yanfang Ye, Nitesh V. Chawla, and Huan Liu. 2022. [Few-Shot Learning on Graphs](#).
- Qiang Zhang, Hongbin Huang, Shangsong Liang, Zaiqiao Meng, and Emine Yilmaz. 2021. [Learning to Detect Few-Shot-Few-Clue Misinformation](#).
- Tong Zhang, Di Wang, Huanhuan Chen, Zhiwei Zeng, Wei Guo, Chunyan Miao, and Lizhen Cui. 2020. [BDANN: BERT-Based Domain Adaptation Neural Network for Multi-Modal Fake News Detection](#). In *2020 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8.
- Jiong Zhu, Yujun Yan, Lingxiao Zhao, Mark Heimann, Leman Akoglu, and Danai Koutra. 2020. [Beyond Homophily in Graph Neural Networks: Current Limitations and Effective Designs](#).

Credibility	2017	2020	2022
High	Trump says he'll allow Kennedy assassination files to be released	Pelosi: Proposal on COVID-19 relief is "one step forward, two steps back"	MPs insist fans heading to World Cup must not be priced out of enjoying a beer
Low	Russian Email Uncovered... Reveals What Really Happened at Trump Jr and Russia Mtg	Pence Destroys Biden's Record: He'd Have Killed 2 Million People Fighting COVID	Joe Biden: 'Inflation Is Going to Get Worse' if Republicans Win Despite Core Inflation Rising to 40-year High

Table 5: A series of headlines taken from the NELA-GT corpora. All articles come from around Oct. 15th in their respective year.

Rumour Status	Anchor	Near	Far
True	It's been 10 years since Heath Ledger died of an accidental drug overdose. Since	Who cut the head off of the General Pickens statue and what is going on in the Cooper	When Pretty Little Liars and Teen Wolf collide we get Truth or Dare. Well, not really,
Fake	A man reportedly got his finger bitten off at a Beyoncé concert! The shocking twist: It wasn't	A white witch from North London has urged Hollywood star Angelina Jolie to cease	Is Cher concerned Chaz Bono will die from his weight issues? That's the claim from

Table 6: A similar set of snippets taken from the GossipCop dataset. The left column provides an anchor document, the middle another document near the anchor, and the right column a document far away from the anchor.

A Motivating Examples

DISCLAIMER: the chosen examples were taken verbatim from various malicious content corpora. They do not reflect the views of the authors.

As established in Section 1, malicious content and its social context evolves. This can happen quickly, and results in text that is very different from already seen forms of malicious content.

Table 5 shows such change, depicting titles from articles published by low (i.e. those that often publish severely biased or false news) and high credibility news sources, as found in the NELA-GT corpora (Horne, 2019; Horne and Gruppi, 2021, 2023).

Current models are adept at filtering out malicious content as in their training datasets, but quickly degrade when presented with novel content. For example, Table 5 shows substantial high-level semantic change across the years. Models relying on surface-level features will fail as new events spawn new content.

Currently, no existing social-network malicious content dataset captures this level of evolution. In fact, existing datasets are completely static, implicitly assuming the full network (users, content, and their connections) will be available at inference

time. This paper argues that this assumption is a significant reason why malicious content continues to propagate unabated, despite the impressive classification scores reported in earlier work.

To showcase this lack of diversity, Table 6 depicts a similar array of texts, sampled from the GossipCop graph. While a variety of topics are discussed, the overarching subject remains the same. Comparatively, relying on surface level features can already lead to strong classification performance. Simply put, in this dataset, models need not account for evolving content.

In lieu of large, temporally diverse graph datasets, we propose an evaluation framework that approximates these effects through generalisation to new graphs. Requiring adaptation from a minimal set of examples, with limited social context will serve as a much better measure of the inference time performance of community models for malicious content detection. There is little point in good within dataset performance, when unseen content forms are free to cause harm.

Put otherwise, when it comes to malicious content detection, we want models that are able to filter out tomorrow's hate speech posts and fake news articles, not those seen yesterday.

B Additional Details on Datasets Used

Rehydrating many years after the datasets were released, not all documents and users could be recovered. This results in empty, missing or isolated documents. The collected graph datasets are thus subgraphs of the one presented in [Shu et al. \(2019a\)](#), [Cui and Lee \(2020\)](#) and [Waseem and Hovy \(2016\)](#). The additional preprocessing steps needed are described here:

1. **Tokenization:** all found documents were collected and tokenized. Any empty documents, or documents yielding only special tokens, were removed
2. **Document-User Interactions:** where possible, user and doc-user interactions were collected. One issue with GossipCop, as identified by [Shu et al.](#), is the inclusion of bots. These ‘users’ tend to disproportionately interact with documents of a single class, both in terms of volume and proportion. Therefore, following the recommendation made by [Chandra et al. \(2020\)](#), users sharing more than 30% of the documents of any class were removed. The type of doc-user interactions in Twitter-HateSpeech differs, resulting in a very small pool of racist users, so this restriction was relaxed. Documents without *any* user interactions, were also removed.
3. **User-user Interactions:** all remaining users and their interaction with other users were parsed at this point. To further reduce the number of bots, the top 1% most active users were removed on GossipCop. Then, again only on GossipCop, to further sparsify the graph, only the top 30k users were kept. Isolated documents were once again removed.

The effect of each filtering step and additional statistics of the dataset graph prior to generating episodic subgraphs, are presented in [Table 7](#).

C Additional Information on Few-shot Subgraph Sampling

[Algorithm 1](#) presents the few-shot subgraph sampling pseudocode. Although the models used can only aggregate information from at most an $r = 2$ radius subgraph, the initial graph can be expanded to larger r . This is to ensure at least k -shot examples of each label is present. For the used graph

$r = 5$ usually contains the vast majority of the graph, and would only be needed for extremely sparse areas. In practise, however, the k -shot was achieved by $r = 3$ in all situations.

The random walk subsampling dramatically reduces the number of nodes and edges present in the subgraph. A similar strategy was employed by GraphSAINT ([Zeng et al., 2020](#)), resulting in efficiency improvements for a variety of inductive graph learners. We set our random walk length to 5 for all experiments, preferring fewer document nodes (with smaller walk length requiring more roots to get to the same node budget). An approximate budget of 2048 was used during training and evaluation.

Important statistic on the produced subgraphs, for both the support and query sets, are presented in [Tables 8 and 9](#).

D Hyperparameters

D.1 GossipCop Training

The set of hyperparameters used for pre-training meta-learners on GossipCop are presented in [Table 10](#). The model size was left constant. Input dimensions were determined by RoBERTa, and set to 768. The GAT attention heads, 3, used an internal dimensionality of 256, and all heads were concatenated afterwards. After GAT processing, representations were fed through a two layer ReLU activated MLP of dimensionality 64, before being classified.

All other hyperparameters were tuned on the validation graphs generated by inductive stratified 5-fold cross validation. Dropout was applied on the internal hidden dimensions. Dropout was applied node-wise on the initial node embeddings, stochastically setting entire nodes to 0. In our case specifically, this essentially means converting documents into users. Attention dropout was applied to the attention weights produced by the GAT layers ([Equation 2](#)). Alternatively, we also experimented with *node masking*, but we did not see a big difference in performance. ([Mishra et al., 2020](#)) AdamW ([Kingma and Ba, 2017](#); [Loshchilov and Hutter, 2019](#)) was the outer-loop optimizer, with only the learning and weight decay terms tuned. The different algorithms required substantially different numbers of gradient update before convergence. Early stopping was used, with patience equal to 10% of the maximum allowed number of steps. Most checkpoints converged well before that

Group	Metric	GossipCop	CoAID	Twitter Hate Speech
Typology	Task	Rumour Verification	Fake News Detection	Hate Speech Classification
	Domain	Celebrity Gossip	COVID-19	Entertainment
	Labels	True Fake 77.12% 22.88%	True Fake 94.72% 5.28%	Racism Sexism None 11.97% 19.43% 68.60%
	Doc-user Interactions	Retweet	Retweet	Authorship
Length	Mean	352.99	71.34	24.42
	Std. Dev	165.4	42.44	9.38
	Median	405	93	25
Missing Documents	Not Found	1168	0	0
	Empty	488	0	0
Users	#Users (pre filter)	549225	5524	1875
	Unique Users - 0	215072	5062	5
	Unique Users - 1	384760	462	527
	Unique Users - 2	N/A	N/A	1648
	Too active	36	0	0
User-Doc Interaction	Isolated Docs-0	1261	3635	0
	Isolated Docs-1	224	875	0
	Isolated Docs-2	N/A	N/A	0
	Mean	2.3	1.06	8.64
	Std. Dev	31.27	0.41	147.49
	Median	1	1	1
	$\mathbb{E}[\log(x)]$	0.25	0.03	0.49
	Geom. Mean	1.29	1.04	1.63
User Truncation	Most Active	5213	0	0
	Least Active	486087	0	0
	# Doc. Incident	27148	4284	1875
	# Doc. Non-incident	2081	0	0
User Degrees	Mean	3445.78	2313.63	16.71
	Std. Dev	1553.34	2373.26	149.15
	Median	3199	1281	4
	$\mathbb{E}[\log(x)]$	8.03	6.95	1.54
	Geom. Mean	3070.33	1042.83	4.65
Graph	#Nodes	46846	5006	18076
	User-doc Edges	284757	4520	16201
	User-user Edges	859097	23604	7561
	Total Edges (uni)	2334554	61254	65600
	Density	2.13E-03	4.89E-03	4.00E-04

Table 7: Dataset filtering and additional statistics on documents and interactions.

Dataset	Metric	Support				
		Mean	Std. Dev.	Q25	Median	Q75
GossipCop	#Nodes	1994.05	242.98	2049	2050	2051
	#Edges	111293.61	41409.85	84358.75	115236	130180.25
	#Docs.	200.62	232.94	40.25	92	256
	Density	0.06	0.02	0.04	0.06	0.06
	Prop. Docs.	9.90%	11.35%	2.15%	4.59%	12.58%
	Deg. Cent.	1.08E-02	6.53E-03	7.39E-03	9.51E-03	1.25E-02
	Eigen Cent.	2.40E-03	2.44E-03	1.65E-03	2.10E-03	2.75E-03
CoAID	#Nodes	1796.66	448.43	1765.25	2048	2049
	#Edges	30464.29	6866.59	27552	33672.5	35229.25
	#Docs.	208.55	107.16	114	243	302
	Density	0.02	0.03	0.02	0.02	0.02
	Prop. Docs.	10.88%	4.62%	7.25%	12.11%	14.80%
	Deg. Cent.	7.46E-03	9.75E-03	3.42E-03	4.01E-03	6.34E-03
	Eigen Cent.	3.18E-03	3.88E-03	1.09E-03	1.86E-03	3.18E-03
TwitterHS	#Nodes	2050.55	2.45	2050	2051	2051
	#Edges	10381.76	515.17	10070.25	10312	10564.25
	#Docs.	1716.29	27.05	1707.25	1721	1732
	Density	0	0	0	0	0.01
	Prop. Docs.	83.70%	1.33%	83.26%	83.91%	84.45%
	Deg. Cent.	1.46E-03	1.88E-06	1.46E-03	1.46E-03	1.46E-03
	Eigen Cent.	2.48E-04	1.20E-03	1.40E-04	1.50E-04	1.62E-04

Table 8: Additional statistics on the subgraphs generated by the proposed sampling procedure on the support graphs. Each row header gives the dataset. The metrics provided include: the number of nodes, number of edges, number of document nodes, the graph density, the proportion of document nodes, the degree centrality of document nodes, and the eigen centrality of document nodes.

Dataset	Metric	Query				
		Mean	Std. Dev.	Q25	Median	Q75
GossipCop	#Nodes	3046.03	2821.88	609.75	2697.5	4779.25
	#Edges	75547.65	187261.29	5549	18839	48589
	#Docs.	2181.1	1976.14	281	1616	4108.5
	Density	0.04	0.09	0	0.01	0.04
	Prop. Docs.	68.76%	31.75%	43.82%	85.18%	94.75%
	Deg. Cent.	4.32E-04	8.01E-04	1.32E-04	2.37E-04	4.48E-04
	Eigen Cent.	1.50E-04	3.78E-04	2.66E-05	6.56E-05	1.40E-04
CoAID	#Nodes	48.07	94.01	7	18	44
	#Edges	525.98	1717.85	24	80.5	257.75
	#Docs.	1.86	1.58	1	1	2
	Density	0.74	0.6	0.27	0.56	1.07
	Prop. Docs.	13.93%	14.36%	2.93%	8.33%	20.00%
	Deg. Cent.	1.67E-03	2.73E-03	6.02E-04	8.02E-04	1.60E-03
	Eigen Cent.	5.47E-04	4.44E-03	1.44E-05	4.83E-05	1.32E-04
TwitterHS	#Nodes	1485.73	1009.45	220	2056	2571
	#Edges	4511.72	2967.31	1774	6168	7711
	#Docs.	1476.97	1017.5	22	2053	2569
	Density	0.23	0.49	0	0	0.13
	Prop. Docs.	81.80%	32.20%	81.25%	99.85%	99.92%
	Deg. Cent.	2.97E-04	1.08E-19	2.97E-04	2.97E-04	2.97E-04
	Eigen Cent.	4.12E-03	6.36E-03	7.01E-08	3.70E-07	1.39E-02

Table 9: Same as Table 8, but now for the query graphs.

Parameter	Full	Subgraphs	MAML-LH	MAML-RH	ProtoNet	ProtoMAML
GAT Hidden Dim				256		
GAT Heads				3		
CLF Dim				64		
Training & Adaptation						
Dropout	0.5	0.4	0.5	0.5	0.5	0.5
Node Dropout	0.1	0.1	0.1	0.1	0.1	0.1
Attn. Dropout	0.1	0.0	0.1	0.1	0.1	0.1
LR	2.50E-03	2.50E-03	5.00E-04	1.00E-03	1.00E-03	1.00E-03
Weight Decay	1.00E-02	5.00E-02	5.00E-02	5.00E-02	5.00E-02	5.00E-02
Batch size	N/A	32	N/A	N/A	N/A	N/A
Updates	100	300	2560	2560	2560	2560
Decay Updates	5	15	128	128	128	128
Decay Factor				0.7943		
Patience	10	30	256	256	256	256
Inner Loop Adaptation - Training						
LR - GAT	N/A	N/A	1.00E-03	1.00E-02	N/A	1.00E-02
LR - CLF Head	N/A	N/A	1.00E-03	5.00E-02	N/A	5.00E-02
T_{inner}	N/A	N/A	1	5	N/A	10
Inner Loop Adaptation - High Adaptation Evaluation						
LR - GAT	N/A	1.00E-02	5.00E-03	5.00E-02	N/A	1.00E-02
LR - CLF Head	N/A	5.00E-01	5.00E-02	5.00E-02	N/A	5.00E-01
T_{inner}	N/A	25	25	25	N/A	25

Table 10: Hyperparameters used for pre-training models on GossipCop, and the high-adaptation evaluation during the transfer experiments.

point. The learning rate was decayed in a step-wise manner, every 5% of the maximum number of steps, with a minimum learning rate at 0.01 of the initial value.

The inner-loop learning rate saw more variation between the different learning algorithms. MAML-LH performs best under minimal adaptation, yielding a single step inner loop with a low learning rate. Resetting the head in each episode, instead, forces adaptation, reflected in a larger, more aggressive inner loop. PROTOMAML, finally, reaches minimum validation loss only with large amounts of inner-loop adaptation.

During generalisation to unseen graphs, a more aggressive adaptation strategy was applied to most models. This was not tuned on the test set. Instead, the highest values possible were, such that no infinities appeared in the output logits.

D.2 Additional Training Details

All experiments were conducted on a Linux-based SLURM-based academic cluster. Nodes consisted

of an Intel Xeon Platinum 8360Y CPU with 18 cores in user at 2.4 GHz, a single NVIDIA A100 GPU accelerator (yielding 40 GiB of HMB2 memory) and 128 GiB of DDR4 memory. The code is written exclusively using Python 3.10.6, PyTorch 1.13.0, built with CUDA 11.7. Graph modelling utilized PyTorch Geometric 2.3.0. All experiments were conducted under random seed 942. For local development we use Ubuntu 20.04.6 LTS (GNU/Linux 5.15.90.1-microsoft-standard-WSL2 x86_64).

E Comparing MAML-LH & MAML-RH

As originally described in Finn et al. (2017), MAML resets its classification head each episode in order to adapt to a new task with new labels. However, in our case, pre-training MAML uses only a single task, with a fixed label definition. Therefore, the classification head can be learned in the outer loop along with the other meta-initialized models. This setting we dubbed MAML-LH (learned head), and has the benefit of requiring

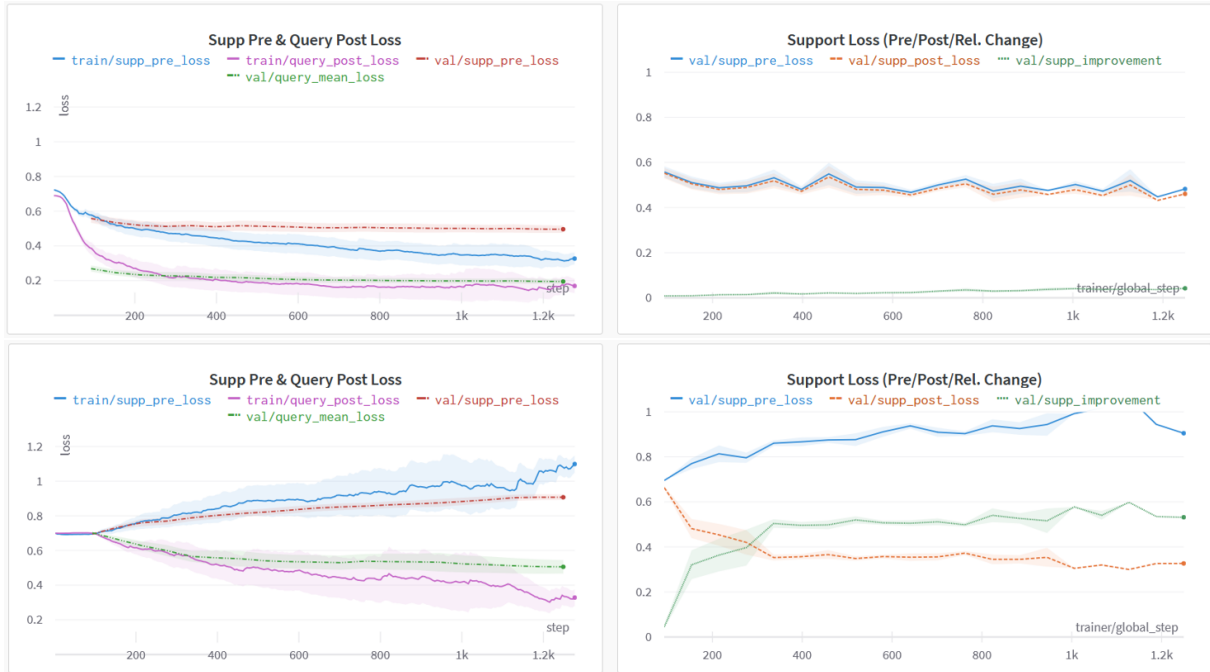


Figure 3: GossipCop training losses for (t) MAML-LH and (b) MAML-RH. The left column gives loss of the models on the support and query sets. Support loss is computed prior to the first adaptation step (blue and dashed orange lines), query loss after the last adaptation step (pink and dashed green lines). The left column provides the support loss prior to the first adaptation step (blue line) and after the last adaptation step (orange line). Finally, the green line gives the relative improvement of the support set loss.

less inner-loop adaptation; at least the classifier head is already task-specific. The standard MAML setup we dubbed MAML-RH (random head).

On GossipCop, at least, this had a dramatic effect on the degree of adaptation, and as a result, on the final performance scores. This can be seen in Figure 3, with the top row of figures giving various losses for MAML-LH, and the bottom row for MAML-RH. The left column of figures present the loss on the support set (prior to adaptation), along with the loss on the query set (post adaptation), on both the train and validation splits. MAML-LH acts like a standard machine learning model. After random performance initially, train loss decreases steadily on both graphs, whereas validation loss stagnates earlier on. The query loss is lower, but as we’re using foMAML with a disjoint support–query split this is to be expected (the model is never directly optimised on the support graphs). MAML-RH, on the other hand, shows rapid divergence in the support loss, while the support loss decreases as usual.

These loss patterns indicate a distinction between the two operating modes of MAML trained models. MAML-LH, seeing a stable learning objective, learns to initialise using representations

optimal to all tasks. In the meta-learning literature this corresponds to ‘feature reuse’, and makes MAML-LH similar to the ANIL (Raghu et al., 2020) MAML variant. MAML-RH, on the other hand, has to leverage the support set to rapidly adapt in order to achieve non-random performance on the query set. Its initial weights are not usable for representation learning, but rather for optimizing itself into a representation learner. This phenomenon is called ‘rapid learning’. This makes MAML-RH more reminiscent of ‘true’ MAML, or the BOIL variant (Oh et al., 2021).

This is made more clear in the right column of Figure 3. Here the support loss before and after adaptation is shown, with a green line also indicating the relative decrease. For MAML-LH, there is barely any difference between the two, but the overall line is already relatively low; it simply does not need to adapt to achieve generalisable representations. MAML-RH, again, is a polar opposite, with diverging initial support loss, but a much lower final loss. The relative improvement is indicative of a model that ‘learns-to-learn’.

In order to test which meta-learning property is more important for the task at hand, both were trained and applied. While MAML-LH clearly

showed itself superior on GossipCop, we initially thought that added bias toward rapid adaptation might aid generalisation with MAML-RH.

F Homophily

Graph homophily is the property of nodes to prefer attaching to similar nodes. In social network users, where interaction usually denotes some form of kinship, this follows naturally from people’s social relationships. In malicious content detection, this is likely a relevant feature as well. Our perception of real and fake news is influenced by our social network neighbourhood, and propagation usually occurs in homophilic settings (Sun et al., 2022).

The effect of homophily on GNN performance remains an open question. In a homophilic setting, node representations will be built from nodes of the same class, whereas in heterophilic settings, node representations contains representations of nodes of different classes. A third setting, less explored, is the case of randomness: nodes are just as likely to attach to nodes of a different class as its own. Zhu et al. (2020) define a homophily metric that measures the global propensity of links between similar nodes. They show that GNNs can fail in heterophilic settings, with an MLP being more effective, despite the graph setting. For a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, they define homophily as,

$$h^{(\text{edge})} = \frac{|\{(u, v) | (u, v) \in \mathcal{E} \vee y_u = y_v\}|}{|\mathcal{E}|}, \quad (8)$$

i.e., the ratio of edges from nodes to similarly labelled nodes, to all edges.

Later papers dispute the claim that GNNs cannot perform under heterophily. Ma et al. (2021) find that GNNs require certain conditions to be met for class separation (used as a proxy for classification performance). Specifically, they indicate that as long as nodes of the same label share similar neighbourhood patterns, node representations will become more similar, despite dissimilar neighbouring nodes.

Lim et al. (2021) take issue with the definition of the homophily metric. Graphs with many node labels will naturally be less homophilic. They propose a metric that measures homophily while correcting for a randomly connected null model where

nodes. Extending to neighbourhoods of radius r ,

$$h_r^{(\text{class insen.})} = \frac{1}{|C| - 1} \sum_{c=1}^{|C|} [h_c^{(\text{neigh.})} - p_c]_+, \quad (9)$$

$$h_{r,c}^{(\text{neigh.})} = \frac{\sum_{v \in \mathcal{V}_c} |\{u | \mathcal{N}_r(v) \wedge y_u = y_v\}|}{\sum_{v \in \mathcal{V}_c} |\{u | \mathcal{N}_r(v)\}|},$$

$$p_c = \frac{|\mathcal{V}_c|}{|\mathcal{V}|}.$$

It may be interpreted as measuring the expected *excess* homophily present in neighbourhoods about nodes of class c .

The proposed metrics for measuring homophily summarize whole graphs, make no distinction between user or document nodes, and do differentiate between a randomly connected graph or a heterophilic graph. In the proposed method of this paper, a homophily metric must be comparable across many subgraphs. As such, inspired by the measure of assortativity, introduced by Newman (2003), we slightly modify the homophily definition as,

$$\hat{h}_{r,c}^{(\text{subgraph})} = \frac{1}{|\mathcal{V}_c|} \sum_{v \in \mathcal{V}_c} \frac{h_{r,c}^{(\text{neigh.})}(v) - p_c}{1 - p_c}, \quad (10)$$

$$h_{r,c}^{(\text{neigh.})}(v) = \frac{|\{u | \mathcal{N}_r^{(\text{docs})}(v) \wedge y_u = y_v\}|}{|\{u | \mathcal{N}_r^{(\text{docs})}(v)\}|}.$$

For a subgraph, it defines the homophily of class c as the expected ratio of homophilic nodes in the r radius neighbourhood of nodes $v \in \mathcal{V}_c$, in excess of a random graph. The use of a neighbourhood to compute $h_{r,c}^{(\text{neigh.})}$ is deliberate. It can now measure the effect of other document nodes on the representation of the centre node. The division by $1 - p_c$ normalizes the excess: a score of 1 is achieved only if fully homophilic, 0 if random, and $-\frac{p_c}{1 - p_c}$ if perfectly heterophilic. This allows for interpreting homophily on a scale.

More important, Equation 10 is applicable for support graphs (which have multiple nodes of each label), and query graphs (which have a single node from a single label). For the support graph, nodes’ homophily scores are averaged, per class, to produce a single summary statistic. For the query graphs, the non-labelled nodes scores are simply omitted.

Rather than presenting a single summary metric, the distribution of homophilic nodes can be observed for sampled subgraphs. This is presented

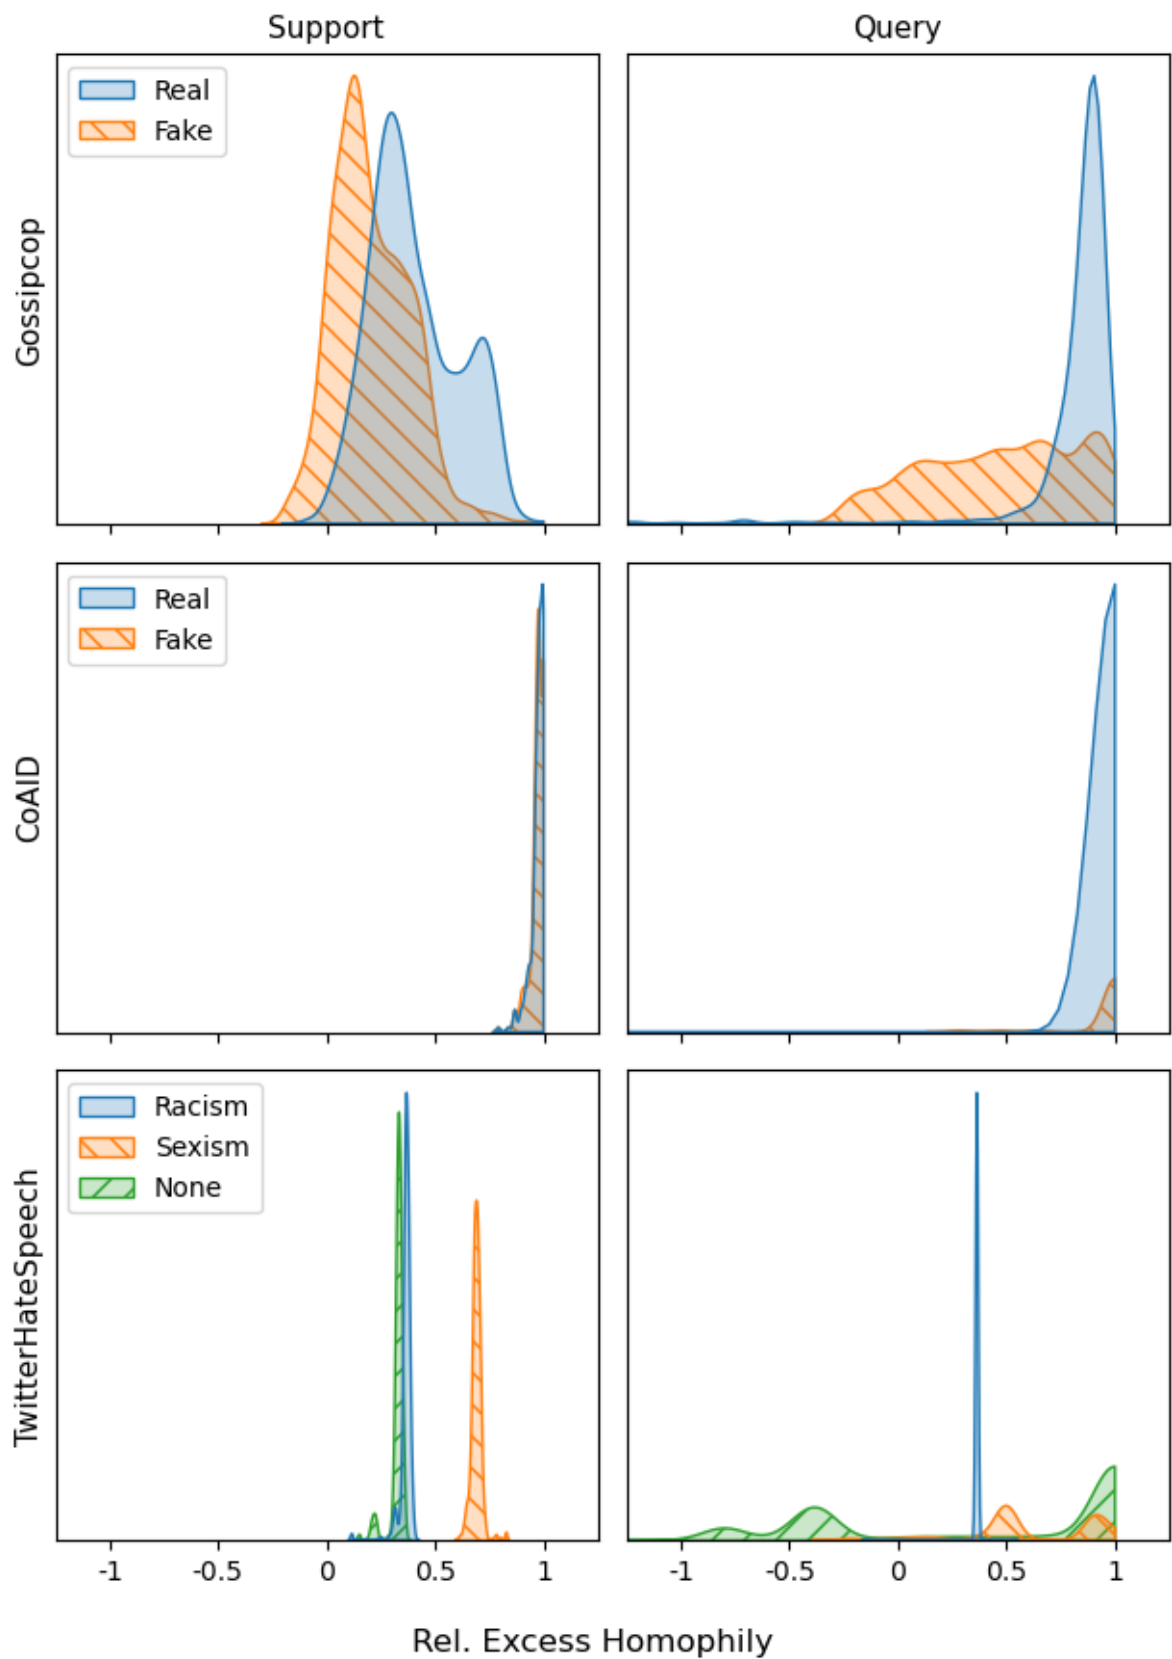


Figure 4: Kernel density estimates for the distribution of relative excess homophily (Equation 10) for sampled subgraphs. The left column present user-centred sampled graphs, the right column gives the r -radius neighbourhoods about document nodes from the query graph. The different rows give different datasets. On the x-axis, 0 corresponds to a random graph, 1 to a perfectly homophilic graph. Values below 0 indicate heterophily.

in Figure 4. The left column presents the distribution of homophily scores for support graphs, right the query graphs. The rows present the different datasets.

For GossipCop, the support graphs show a relatively broad distribution of homophily scores, with modes between the 0-0.5 range. The query graph, however, shows a significant difference between the two classes, with the real documents being highly homophilic, and fake less so. In other words, there are fake documents whose neighbourhood consists primarily of real documents. CoAID shows more consistent behaviour, with both node classes being extremely homophilic.

TwitterHateSpeech is somewhat of an outlier, with substantial differences between the classes. Again, the collection procedure used by Waseem and Hovy led to a small number of extremely active, racist users. Furthermore, user-document links represent authorship, not tweet/re-tweet interactions. As a result, document representations are built up entirely out of the *other* posts by the same user. As a result, the racist class has a narrow distribution, that indicates slight homophily. The innocuous class, ‘None’, gives a bimodal distribution in the query graphs. Many innocuous documents seem to be produced by regular Twitter users, whereas a large portion come from racist and sexist users (a large bump in the heterophilic range). Only the sexist class seems to be consistently homophilic. As a result, this dataset is relatively noisy, with representations being influenced by dissimilar documents.

G Extended Generalisation Results

G.1 CoAID

The results presented in Figure 2 are presented in tabular form in Tables 11 and 12. Since CoAID matches, approximately, the pre-training task used in GossipCop, models were initially adapted using the same classification head (in case of SUBGRAPHS and MAML-LH) and inner-loop learning parameters. This corresponds to the low-adaptation setting, presented in Table 11. This setting estimates direct domain transfer, much like testing the Subgraphs model at $k = 0$. Performance proved disappointing for most models, with the most aggressive adapter (PROTOMAML), clearly exceeding all other tested models.

Therefore, we conducted a second round of experiments with similarly aggressive inner-loop

learning, presented in Table 12. The only models exempted, were SUBGRAPHS at $k = 0$ and the PROTONET, as neither adapts. All models benefited from the more aggressive inner loop, indicating that the generalisation is not trivial. All-in-all, the highest achieved MCC was 0.1709, for PROTOMAML at $k = 8$, corresponding to an F1-Fake of 0.1841. While low relative to other F1 scores reported, this should be compared to a class prevalence of 5%.

G.2 TwitterHateSpeech

Similarly, the TwitterHateSpeech results presented in Figure 2 are presented in tabular form in Table 13. Having learnt from CoAID, only the high-adaptation hyperparameters were used. Lower k -shot values see the ‘None’ class dominate in terms of classification scores. Performance on the minority classes increased with larger k -shot values, at the cost of reduced innocuous tweet performance. Ultimately, ProtoMAML manages this trade-off best, with gradually increasing MCC scores.

G.3 Ablating GossipCop Pre-Training

Tables 14 and 15 show additional results pertaining to the ablation experiment described in Section 5.4. The largest addition, is the inclusion of the other type of GBML algorithm, MAML. The comparison model used is MAML-RH.

On CoAID, MAML-RESET yields ‘always positive’ models, giving constant MCC scores of 0. ProtoMAML, however, proves reasonably robust, with PROTOMAML-RESET performance that exceeds trained MAML-RH. The same effect also holds on TwitterHateSpeech, with PROTOMAML only overcoming its reset counterpart in the larger k -shot settings.

G.4 Extreme k -shot

To test the capacity of the meta-learners, a limited extension of the TwitterHateSpeech experiment was conducted. Instead of limiting ourselves to $k = 16$ examples, we increased to $k = 256$. Only PROTONET was used. Results are depicted graphically in Figure 5 and given numerically in Table 16.

We fully expect diminishing returns. Our graph setting implies that the k labelled nodes are already present in the support graph, just with its label masked. Unmasking additional labels should provide little additional information to the model; a good graph learner can already infer the masked labels. In fact, as discussed in Section 5.3, under

k	Method	F1		MCC
		Real	Fake	
0	SUBGRAPHS	0.2445 (0.2445, 0.2446)	0.1121 (0.112, 0.1122)	0.0398 (0.0394, 0.0402)
	SUBGRAPHS	0.7164 (0.71, 0.7227)	0.1306 (0.129, 0.1321)	0.0642 (0.0628, 0.0656)
	MAML-LH	0.4225 (0.4224, 0.4227)	0.1162 (0.1161, 0.1164)	0.0593 (0.0592, 0.0593)
	MAML-RH	0.8241 (0.8198, 0.8285)	0.1561 (0.1548, 0.1575)	0.1164 (0.1149, 0.1179)
	PROTONET	0.7007 (0.6928, 0.7085)	0.1404 (0.1388, 0.1421)	0.0731 (0.0692, 0.0771)
4	PROTOMAML	0.7867 (0.7791, 0.7942)	0.1767 (0.1746, 0.1788)	0.1321 (0.1279, 0.1363)
	SUBGRAPHS	0.2219 (0.2218, 0.222)	0.1044 (0.1043, 0.1046)	0.0377 (0.0372, 0.0383)
	MAML-LH	0.5693 (0.5691, 0.5695)	0.1096 (0.1094, 0.1098)	0.0571 (0.057, 0.0572)
	MAML-RH	0.8243 (0.8208, 0.8278)	0.1567 (0.1556, 0.1577)	0.1234 (0.1221, 0.1247)
	PROTONET	0.7540 (0.7493, 0.7588)	0.1471 (0.1458, 0.1484)	0.1179 (0.1154, 0.1205)
8	PROTOMAML	0.8300 (0.826, 0.8341)	0.1799 (0.1784, 0.1814)	0.1616 (0.1594, 0.1638)
	SUBGRAPHS	0.2387 (0.2386, 0.2389)	0.0974 (0.0972, 0.0975)	0.0370 (0.0363, 0.0377)
	MAML-LH	0.5740 (0.5738, 0.5741)	0.1028 (0.1025, 0.103)	0.0551 (0.055, 0.0552)
	MAML-RH	0.8355 (0.8328, 0.8383)	0.1466 (0.1457, 0.1475)	0.1154 (0.1142, 0.1167)
	PROTONET	0.7560 (0.7517, 0.7604)	0.1395 (0.1384, 0.1405)	0.1175 (0.1154, 0.1195)
12	PROTOMAML	0.8331 (0.8295, 0.8367)	0.1675 (0.1662, 0.1687)	0.1583 (0.1566, 0.1601)
	SUBGRAPHS	0.2296 (0.2294, 0.2297)	0.0896 (0.0894, 0.0898)	0.0379 (0.0371, 0.0387)
	MAML-LH	0.5871 (0.5869, 0.5873)	0.0956 (0.0954, 0.0959)	0.0528 (0.0527, 0.0529)
	MAML-RH	0.8315 (0.8291, 0.834)	0.1352 (0.1343, 0.1361)	0.1106 (0.1093, 0.112)
	PROTONET	0.7429 (0.7383, 0.7475)	0.1268 (0.1259, 0.1278)	0.1097 (0.1079, 0.1115)
16	PROTOMAML	0.8323 (0.8291, 0.8355)	0.1530 (0.1518, 0.1541)	0.1508 (0.1493, 0.1523)

Table 11: CoAID transfer results under low adaptation hyperparameters.

k	Method	F1		MCC
		Real	Fake	
0	SUBGRAPHS	0.2445 (0.2445, 0.2446)	0.1121 (0.112, 0.1122)	0.0398 (0.0394, 0.0402)
	SUBGRAPHS	0.727 (0.7199, 0.7341)	0.131 (0.1289, 0.1331)	0.0727 (0.0708, 0.0746)
	MAML-LH	0.7494 (0.7435, 0.7553)	0.1538 (0.1524, 0.1552)	0.1172 (0.1154, 0.119)
	MAML-RH	0.7797 (0.7744, 0.7851)	0.1562 (0.1547, 0.1577)	0.1195 (0.1177, 0.1214)
	PROTONET	0.7007 (0.6928, 0.7085)	0.1404 (0.1388, 0.1421)	0.0731 (0.0692, 0.0771)
4	PROTOMAML	0.7734 (0.7655, 0.7812)	0.1762 (0.174, 0.1784)	0.1383 (0.1343, 0.1422)
	SUBGRAPHS	0.7651 (0.7597, 0.7705)	0.129 (0.1268, 0.1312)	0.0827 (0.0807, 0.0846)
	MAML-LH	0.7679 (0.7633, 0.7724)	0.1525 (0.1512, 0.1537)	0.1232 (0.1215, 0.1248)
	MAML-RH	0.8021 (0.7983, 0.8059)	0.1581 (0.157, 0.1593)	0.1312 (0.1297, 0.1326)
	PROTONET	0.754 (0.7493, 0.7588)	0.1471 (0.1458, 0.1484)	0.1179 (0.1154, 0.1205)
8	PROTOMAML	0.8245 (0.8202, 0.8288)	0.1841 (0.1824, 0.1858)	0.1709 (0.1685, 0.1733)
	SUBGRAPHS	0.7783 (0.774, 0.7825)	0.1316 (0.1297, 0.1334)	0.0921 (0.0904, 0.0939)
	MAML-LH	0.7893 (0.7856, 0.7931)	0.147 (0.1459, 0.1482)	0.1263 (0.1248, 0.1279)
	MAML-RH	0.8181 (0.8149, 0.8213)	0.1505 (0.1495, 0.1516)	0.1281 (0.1267, 0.1295)
	PROTONET	0.756 (0.7517, 0.7604)	0.1395 (0.1384, 0.1405)	0.1175 (0.1154, 0.1195)
12	PROTOMAML	0.8294 (0.8255, 0.8333)	0.1732 (0.1718, 0.1746)	0.1689 (0.1669, 0.1708)
	SUBGRAPHS	0.7922 (0.7887, 0.7956)	0.1314 (0.1298, 0.133)	0.1028 (0.1011, 0.1045)
	MAML-LH	0.806 (0.8033, 0.8087)	0.1385 (0.1375, 0.1396)	0.1277 (0.1263, 0.1291)
	MAML-RH	0.8165 (0.8139, 0.8192)	0.1405 (0.1396, 0.1414)	0.1234 (0.1221, 0.1247)
	PROTONET	0.7429 (0.7383, 0.7475)	0.1268 (0.1259, 0.1278)	0.1097 (0.1079, 0.1115)
16	PROTOMAML	0.8321 (0.8288, 0.8354)	0.1599 (0.1587, 0.1612)	0.1646 (0.1631, 0.1662)

Table 12: CoAID transfer results under high adaptation hyperparameters, with the exception for SUBGRAPHS at $k = 0$ and PROTONET, neither of which adapts during evaluation.

k	Method	F1			MCC
		Racism	Sexism	None	
4	SUBGRAPHS	0.1615 (0.1563, 0.1666)	0.1950 (0.1894, 0.2007)	0.3745 (0.3634, 0.3855)	0.0334 (0.0305, 0.0363)
	MAML-LH	0.1696 (0.1633, 0.1759)	0.2287 (0.223, 0.2344)	0.3554 (0.3461, 0.3646)	0.0580 (0.0539, 0.0621)
	MAML-RH	0.1741 (0.1682, 0.1801)	0.2402 (0.2352, 0.2453)	0.3420 (0.3322, 0.3519)	0.0543 (0.0507, 0.0579)
	PROTONET	0.1949 (0.188, 0.2019)	0.2355 (0.2299, 0.2411)	0.3930 (0.3851, 0.401)	0.0784 (0.0742, 0.0826)
	PROTOMAML	0.1763 (0.1695, 0.1831)	0.2181 (0.2124, 0.2238)	0.3585 (0.3494, 0.3676)	0.0607 (0.0565, 0.0648)
8	SUBGRAPHS	0.1545 (0.1493, 0.1598)	0.2161 (0.2108, 0.2214)	0.3750 (0.3646, 0.3854)	0.0333 (0.0304, 0.0361)
	MAML-LH	0.1802 (0.1736, 0.1869)	0.2552 (0.2499, 0.2606)	0.3340 (0.3248, 0.3432)	0.0817 (0.0781, 0.0853)
	MAML-RH	0.1669 (0.1608, 0.173)	0.2614 (0.2574, 0.2653)	0.3480 (0.338, 0.3581)	0.0515 (0.0483, 0.0546)
	PROTONET	0.2157 (0.2092, 0.2222)	0.2221 (0.2162, 0.2281)	0.4065 (0.3991, 0.4138)	0.0904 (0.0862, 0.0946)
	PROTOMAML	0.1934 (0.1866, 0.2002)	0.2148 (0.2092, 0.2205)	0.3530 (0.3439, 0.362)	0.0699 (0.0657, 0.074)
12	SUBGRAPHS	0.1895 (0.1838, 0.1952)	0.2169 (0.2113, 0.2224)	0.3686 (0.3576, 0.3795)	0.0515 (0.0482, 0.0549)
	MAML-LH	0.1898 (0.183, 0.1966)	0.2543 (0.2487, 0.26)	0.3068 (0.2976, 0.316)	0.0770 (0.0733, 0.0807)
	MAML-RH	0.2286 (0.2235, 0.2336)	0.2525 (0.2468, 0.2581)	0.3185 (0.3084, 0.3285)	0.0791 (0.0757, 0.0824)
	PROTONET	0.2798 (0.2743, 0.2852)	0.2612 (0.2554, 0.267)	0.3803 (0.3731, 0.3876)	0.1134 (0.1092, 0.1177)
	PROTOMAML	0.2545 (0.2475, 0.2616)	0.2554 (0.2492, 0.2616)	0.3503 (0.3412, 0.3594)	0.1109 (0.1066, 0.1152)
16	SUBGRAPHS	0.1996 (0.194, 0.2052)	0.2261 (0.2206, 0.2316)	0.3602 (0.3492, 0.3712)	0.0574 (0.054, 0.0608)
	MAML-LH	0.1874 (0.1805, 0.1942)	0.2595 (0.2538, 0.2652)	0.3072 (0.2978, 0.3165)	0.0778 (0.0741, 0.0815)
	MAML-RH	0.2319 (0.227, 0.2368)	0.2510 (0.2453, 0.2567)	0.3159 (0.3059, 0.3259)	0.0811 (0.0778, 0.0844)
	PROTONET	0.2883 (0.2838, 0.2927)	0.2579 (0.2523, 0.2635)	0.3740 (0.3664, 0.3816)	0.1128 (0.1087, 0.117)
	PROTOMAML	0.3021 (0.2961, 0.308)	0.3077 (0.2999, 0.3155)	0.3163 (0.3051, 0.3276)	0.1354 (0.1303, 0.1404)

Table 13: TwitterHateSpeech transfer results under high adaptation hyperparameters.

k	Method	F1		MCC
		Real	Fake	
4	MAML-RESET	0.971 (0.971, 0.9711)	0 (0, 0)	0 (0, 0)
	PROTOMAML-RESET	0.7869 (0.7674, 0.8065)	0.1716 (0.1673, 0.1759)	0.1191 (0.1114, 0.1268)
8	MAML-RESET	0.9732 (0.9731, 0.9732)	0 (0, 0)	0 (0, 0)
	PROTOMAML-RESET	0.8483 (0.8394, 0.8573)	0.1769 (0.1738, 0.18)	0.1398 (0.1361, 0.1434)
12	MAML-RESET	0.9752 (0.9751, 0.9753)	0 (0, 0)	0 (0, 0)
	PROTOMAML-RESET	0.8470 (0.8368, 0.8572)	0.1652 (0.1621, 0.1683)	0.1304 (0.1262, 0.1347)
16	MAML-RESET	0.9774 (0.9773, 0.9775)	0 (0, 0)	0 (0, 0)
	PROTOMAML-RESET	0.8542 (0.8479, 0.8605)	0.1504 (0.1475, 0.1533)	0.1212 (0.1179, 0.1245)

Table 14: Models ‘transferred’ to CoAID after reset.

k	Method	F1			MCC
		Racism	Sexism	None	
4	MAML-RESET	0.1699 (0.163, 0.1768)	0.1918 (0.1853, 0.1983)	0.3433 (0.3348, 0.3517)	0.0726 (0.0687, 0.0765)
	PROTOMAML-RESET	0.1799 (0.173, 0.1868)	0.1906 (0.1845, 0.1966)	0.3225 (0.3136, 0.3314)	0.0767 (0.0729, 0.0806)
8	MAML-RESET	0.2239 (0.2178, 0.2301)	0.1691 (0.164, 0.1742)	0.3277 (0.3193, 0.3361)	0.0811 (0.0777, 0.0846)
	PROTOMAML-RESET	0.1999 (0.1934, 0.2065)	0.1617 (0.1556, 0.1677)	0.3416 (0.3336, 0.3495)	0.0868 (0.0832, 0.0904)
12	MAML-RESET	0.2315 (0.2264, 0.2366)	0.1348 (0.1292, 0.1404)	0.3125 (0.3056, 0.3194)	0.0969 (0.0938, 0.1001)
	PROTOMAML-RESET	0.2471 (0.2421, 0.2521)	0.1399 (0.1348, 0.1449)	0.2882 (0.28, 0.2965)	0.1025 (0.0997, 0.1053)
16	MAML-RESET	0.2365 (0.2319, 0.241)	0.1294 (0.1239, 0.1348)	0.3065 (0.3003, 0.3127)	0.098 (0.095, 0.1009)
	PROTOMAML-RESET	0.2524 (0.248, 0.2569)	0.1233 (0.1184, 0.1281)	0.287 (0.2789, 0.2952)	0.1052 (0.1025, 0.1078)

Table 15: Models ‘transferred’ to TwitterHateSpeech after reset.

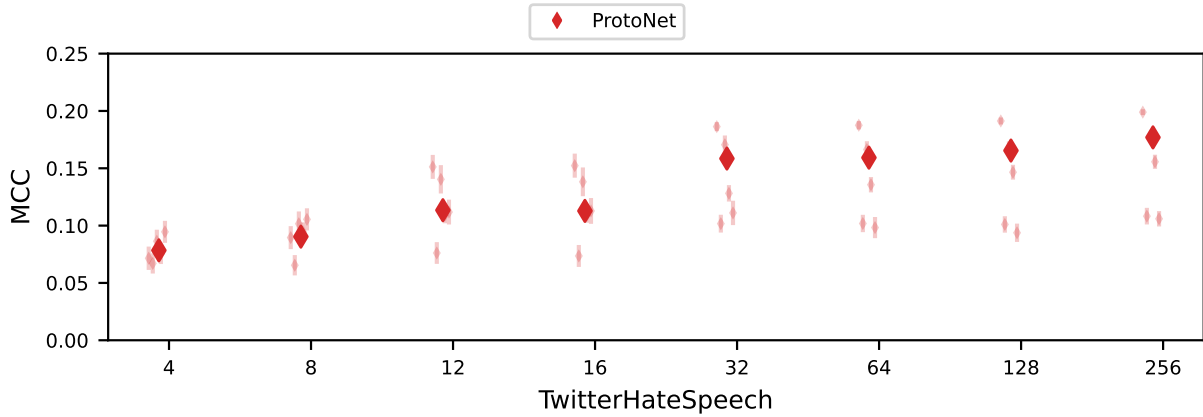


Figure 5: TwitterHateSpeech results using only PROTONET at much larger values of k .

k	F1			MCC
	Racism	Sexism	None	
4	0.1949 (0.1880, 0.2019)	0.2355 (0.2299, 0.2411)	0.3930 (0.3851, 0.401)	0.0784 (0.0742, 0.0826)
8	0.2157 (0.2092, 0.2222)	0.2221 (0.2162, 0.2281)	0.4065 (0.3991, 0.4138)	0.0904 (0.0862, 0.0946)
12	0.2798 (0.2743, 0.2852)	0.2612 (0.2554, 0.267)	0.3803 (0.3731, 0.3876)	0.1134 (0.1092, 0.1177)
16	0.2883 (0.2838, 0.2927)	0.2579 (0.2523, 0.2635)	0.3740 (0.3664, 0.3816)	0.1128 (0.1087, 0.117)
32	0.3374 (0.3365, 0.3382)	0.3126 (0.3099, 0.3153)	0.2947 (0.2902, 0.2993)	0.1585 (0.1557, 0.1613)
64	0.2995 (0.2992, 0.2997)	0.3092 (0.3068, 0.3115)	0.2798 (0.2757, 0.2838)	0.1593 (0.1567, 0.1619)
128	0.2971 (0.2970, 0.2973)	0.3239 (0.3219, 0.3259)	0.2740 (0.2704, 0.2775)	0.1655 (0.1632, 0.1677)
256	0.3001 (0.2999, 0.3003)	0.3432 (0.3416, 0.3449)	0.2713 (0.2680, 0.2746)	0.1770 (0.1750, 0.1789)

Table 16: TwitterHateSpeech transfer results using only PROTONET at much larger values of k .

heterophily, one might expect the additional labels to provide additional noise for the innocuous class.

Precisely this can be observed in Table 16. While the MCC score does increase steadily, it comes at the cost of reduced F1 in the ‘None’ class, and stagnation in the ‘Racism’ class. The only class that sees improvements at very high k -shot values, is the homophobic ‘Sexist’ class.