

---

# Uncertainty for Active Learning on Graphs

---

**Dominik Fuchsgruber\***

Technical University of Munich  
d.fuchsgruber@tum.de

**Tom Wollschläger\***

Technical University of Munich  
t.wollschlaeger@tum.de

**Bertrand Charpentier**

Technical University of Munich  
b.charpentier@tum.de

**Antonio Oroz**

Technical University of Munich  
a.oroz@tum.de

**Stephan Günnemann**

Technical University of Munich  
s.guennemann@tum.de

## Abstract

Uncertainty Sampling is an Active Learning strategy that aims to improve the data efficiency of machine learning models by iteratively acquiring labels of data points with the highest uncertainty. While it has proven effective for independent data its applicability to graphs remains under-explored. We propose the first extensive study of Uncertainty Sampling for node classification: **(1)** We benchmark Uncertainty Sampling beyond predictive uncertainty and highlight a significant performance gap to other Active Learning strategies. **(2)** We develop ground-truth Bayesian uncertainty estimates in terms of the data generating process and prove their effectiveness in guiding Uncertainty Sampling toward optimal queries. We confirm our results on synthetic data and design an approximate approach that consistently outperforms other uncertainty estimators on real datasets. **(3)** Based on this analysis, we relate pitfalls in modeling uncertainty to existing methods. Our analysis enables and informs the development of principled uncertainty estimation on graphs.

## 1 Introduction

Applications in machine learning are often limited by their data efficiency. Active Learning (AL) allows the learner to query an oracle to label specific data points considered *informative*. A popular strategy for query selection is Uncertainty Sampling (US) [1, 2] that relies on uncertainty estimates to measure the informativeness of each data point. Typically, one disentangles the (irreducible) *aleatoric* uncertainty and the (reducible) *epistemic* uncertainty which the *total* uncertainty about a prediction is composed of [3, 4]. This disentanglement is particularly important for AL: The learner might not benefit much from labeling instances with high irreducible uncertainty while acquiring knowledge about data points for which uncertainty stems from reducible sources can be highly informative [5].

For independent and identically distributed (i.i.d.) data, disentangled US methods for AL have demonstrated high data efficiency in benchmarks [1, 2, 5–8]. Existing efforts in AL for interdependent data like graphs neglect the compositional nature of uncertainty [9–13] and limit themselves to a measure of total uncertainty. This leaves it unclear (i) to which extent uncertainty estimators can effectively inform US for graph data, and (ii) whether disentangling aleatoric and epistemic uncertainty has similar benefits to AL as in i.i.d. settings. The complex nature of graph data makes investigating US methods for AL particularly challenging. Uncertainty estimates should not only capture information about node features independently but also model information about their relationships [14].

In this work, we benchmark state-of-the-art uncertainty estimators (see Related Work in Appendix C) against traditional AL strategies and find that US (as well as many other approaches) fall short of surpassing random sampling. We formally approach the question of whether US methods for graphs align well with the AL objective: We derive novel ground-truth uncertainty measures from the

---

\*Equal contribution.

underlying data-generating process and disentangle its aleatoric and epistemic components. This generative perspective enables a principled perspective on handling instance interdependence in uncertainty estimation that we leverage for AL: We prove that querying epistemically uncertain nodes is equivalent to maximizing the relative gain in confidence a Bayesian classifier puts on correctly predicting all unlabeled nodes.

The main **contributions** of our work are:

- We provide the first extensive AL benchmark<sup>2</sup> for node classification including both state-of-the-art uncertainty estimators for US and traditional acquisition strategies. Many of the latter as well as most uncertainty estimators do not outperform random acquisition.
- We derive ground-truth aleatoric and epistemic uncertainty for a Bayesian classifier and formally prove the alignment of US with AL.
- We empirically confirm the efficacy of epistemic US both on synthetic and real data by employing an approximation to our proposed ground-truth uncertainty that outperforms other uncertainty estimators off-the-shelf. This enables future uncertainty estimators to achieve competitive AL performance by building on the principles of our work.

## 2 Ground-Truth Uncertainty from the Data Generating Process

As we show in Appendix D, in contrast to i.i.d. domains, neither traditional nor uncertainty-based AL strategies consistently outperform random queries. This motivates a thorough theoretical examination of whether uncertainty sampling and AL align well for non-i.i.d. problems. However, evaluating the quality of uncertainty estimates is inherently difficult as generally, ground-truth values are unavailable for both the overall predictive uncertainty  $u^{\text{total}}$  and its constituents  $u^{\text{alea}}, u^{\text{epi}}$ . Additionally, since epistemic uncertainty pertains to the knowledge of the classifier, it cannot be defined in a model-agnostic manner. Therefore, we analyze uncertainty from the perspective of the underlying (potentially unknown) data-generating process  $p(\mathbf{X}, \mathbf{A}, \mathbf{y})$  with respect to a Bayesian classifier. This lends itself to a definition of ground-truth uncertainty. In the following, we propose confidence measures and relate them to uncertainty as their inverse  $\text{conf} := u^{-1}$ .

**Definition 2.1.** We define the parametrized Bayesian classifier  $f_{\theta}^*(\mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}}^{\text{gt}})$  in terms of the data generating process  $p(\mathbf{A}, \mathbf{X}, \mathbf{y})$  as the prediction  $\mathbf{c} \in [C]^{|\mathcal{U}|}$  that maximizes:

$$\mathbb{E}_{p(\theta|\mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}}^{\text{gt}})} [\mathbb{P}[\mathbf{y}_{\mathcal{U}} = \mathbf{c} \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}} = \mathbf{y}_{\mathcal{O}}^{\text{gt}}, \theta]] \quad (1)$$

Here, we denote with  $\mathbf{y}_{\mathcal{O}}^{\text{gt}}$  the labels of already observed instances. The predictive distribution  $p(\mathbf{y}_{\mathcal{U}} \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}}^{\text{gt}})$  encapsulates the total confidence of  $f_{\theta}^*$ . The classifier averages its prediction according to a learnable posterior distribution  $p(\theta \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}})$  over its parameters, e.g. weights of a GNN. Marginalization yields the total confidence  $\text{conf}^{\text{total}}$ .

**Definition 2.2.** We define the total confidence  $\text{conf}^{\text{total}}(i, c)$  of  $f_{\theta}^*$  in predicting  $\mathbf{y}_i = c$  as:

$$\mathbb{E}_{p(\theta|\mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}}^{\text{gt}})} [\mathbb{P}[\mathbf{y}_i = c \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}} = \mathbf{y}_{\mathcal{O}}^{\text{gt}}, \theta]] \quad (2)$$

Intuitively, the total confidence captures aleatoric factors through the inherent randomness in the data generating process  $p(\mathbf{y}, \mathbf{A}, \mathbf{X})$ . Epistemic uncertainty is incorporated by conditioning on a limited set of observed labels  $\mathbf{y}_{\mathcal{O}}^{\text{gt}}$ . With a growing labeled set irreducible errors will increasingly dominate total predictive uncertainty. In the extreme case where all labels but one have been observed, i.e.  $\mathcal{O} = \mathcal{V} \setminus \{v_i\}$ , remaining uncertainty only stems from aleatoric factors.

**Definition 2.3.** We define the aleatoric confidence  $\text{conf}^{\text{alea}}(i, c)$  of  $f_{\theta}^*$  in predicting  $\mathbf{y}_i = c$  as:

$$\mathbb{E}_{p(\theta|\mathbf{A}, \mathbf{X}, \mathbf{y}_{-i}^{\text{gt}})} [\mathbb{P}[\mathbf{y}_i = c \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{-i} = \mathbf{y}_{-i}^{\text{gt}}, \theta]] \quad (3)$$

Here, we denote with  $\mathbf{y}_{-i} = \mathbf{y}_{-i}^{\text{gt}}$  that all nodes excluding the predicted node  $i$  are observed as their true values. All remaining lack of confidence is deemed irreducible. Lastly, we define epistemic confidence by comparing aleatoric factors to the overall confidence. As both are defined probabilistically, we consider their ratio:

<sup>2</sup>Find our code at <https://figshare.com/s/95911d5944f57d7ab445?file=44372885>

**Definition 2.4.** We define epistemic confidence of  $f_\theta^*$  in predicting label  $\mathbf{y}_i = c$  as:

$$\text{conf}^{\text{epi}}(i, c) := \text{conf}^{\text{alea}}(i, c) / \text{conf}^{\text{total}}(i, c) \quad (4)$$

These definitions directly imply a notion of uncertainty:  $\text{u}^{\text{epi}}(i, c) = \text{u}^{\text{total}}(i, c) / \text{u}^{\text{alea}}(i, c)$ . Epistemic US thus labels node  $i$  when the associated total uncertainty  $\text{u}^{\text{total}}(i, \mathbf{y}_i^{\text{gt}})$  is large compared to its aleatoric uncertainty  $\text{u}^{\text{alea}}(i, \mathbf{y}_i^{\text{gt}})$ . It favors uncertain nodes *when the uncertainty stems from non-aleatoric sources*.

**Remark 2.5.** In most applications, including US, monotonic transformations of an uncertainty estimator do not affect its behaviour. Therefore, uncertainty can equivalently be defined as a difference between log-likelihoods instead of likelihood ratios. This definition recovers the well-established additive nature of uncertainty:  $\log \text{u}^{\text{total}}(i, \mathbf{y}_i^{\text{gt}}) - \log \text{u}^{\text{alea}}(i, \mathbf{y}_i^{\text{gt}}) = \log \text{u}^{\text{epi}}(i, \mathbf{y}_i^{\text{gt}})$ . The same holds for logarithmic confidence definitions. Notably, the core result of our work, which we state next, also holds when defining uncertainty in terms of log-likelihoods.

**Theorem 2.6.** *Epistemic uncertainty  $\text{u}^{\text{epi}}(i, \mathbf{y}_i^{\text{gt}})$  of a node  $i$  is equivalent to the relative gain its acquisition provides to the posterior over the remaining true labels:*

$$\text{u}^{\text{epi}}(i, \mathbf{y}_i^{\text{gt}}) = \frac{\mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}}, \mathbf{y}_i = \mathbf{y}_i^{\text{gt}}]}{\mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}}]}$$


Hence, acquiring the most epistemically uncertain node is an optimal AL strategy for  $f_\theta^*$ .

We provide a proof of Theorem 2.6 in Appendix B. Here, we refer to all unobserved labels excluding  $\mathbf{y}_i$  as  $\mathbf{y}_{\mathcal{U}-i}$ . The ratio that is optimized by epistemic US corresponds to the relative increase in the posterior of the true unobserved labels. That is, it compares the probability the classifier  $f_\theta^*$  assigns to the remaining unobserved true labels after acquiring the ground-truth label  $\mathbf{y}_i^{\text{gt}}$  of node  $i$  as opposed to not acquiring its label. High values indicate that the underlying classifier will be significantly more likely to predict the true labels of the remaining nodes after the corresponding query. Thus, a query that maximizes epistemic uncertainty will push the classifier toward predicting the true labels for all remaining unlabeled nodes. This holds for any Bayesian classifier that specifies a posterior  $p(\theta \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}})$  over the parameters of the generative process. For example, fitting the parameters of a GNN is an instance of this framework, where all probability mass is put on one estimate of  $\theta$ . Approaches like Bayesian GNNs explicitly specify this posterior distribution. However, computing exact disentangled uncertainty requires access to unavailable labels  $\mathbf{y}_{\mathcal{U}}$  and is therefore impractical. Our analysis motivates the development of tractable approximations to these quantities. Novel US approaches can directly benefit from the theoretical optimality guarantees that this work provides.

### 3 Uncertainty Sampling with Ground-Truth Uncertainty

To support these theoretical claims, we employ US using the proposed ground-truth epistemic uncertainty as an acquisition function. The uncertainty measures depend on the data-generating process, which is not known in real-world applications. Furthermore, computing aleatoric uncertainty necessitates access to unavailable labels  $\mathbf{y}^{\text{gt}}$ . We address these issues by proposing a simple approximate disentangled uncertainty estimator that closely follows our analysis. Our results underline the applicability of our theoretical findings to real-world problems. Additionally, we supply an evaluation on synthetic data for which the generating process is known in Appendix H.1. In the following, we outline the concepts behind two different paradigms to approximate ground-truth uncertainty. We provide an extensive formal description of both algorithms in Appendix J.

**Multiple Pseudo-Labels (MP).** According to Equation (4), one way to obtain epistemic confidence is to approximate total and aleatoric confidence separately and take their ratio. Total confidence is defined in terms of the unknown underlying data generating process  $p(\mathbf{A}, \mathbf{X}, \mathbf{y})$  (Definition 2.2). We approximate it directly from the predictive distribution of the classifier  $f_\theta$ , thereby assuming  $f_\theta$



**Figure 1:** Approximation of disentangled uncertainty against baselines on CoraML. Our framework works well both using multiple pseudo-labels (MP) or taking an expectation over each of them (ESP).

to implicitly model the data distribution:

$$\text{conf}^{\text{total}}(i, c) \approx f_{\theta}(\mathbf{A}, \mathbf{X})_{i,c}$$

Aleatoric confidence (Definition 2.3) also conditions the marginal distribution over  $\mathbf{y}_i$  on unavailable labels  $\mathbf{y}_{\mathcal{U}-i}$ . Therefore, we use the predictions of  $f_{\theta}$  as pseudo-labels

$\hat{\mathbf{y}}_i = \text{argmax}_c f_{\theta}(\mathbf{A}, \mathbf{X})_{i,c}$  and temporarily augment the dataset with *multiple pseudo-labels* by setting  $\mathbf{y}_{\mathcal{U}-i} = \hat{\mathbf{y}}_{\mathcal{U}-i}$ . The predictive distribution of an auxiliary classifier  $f_{\hat{\theta}_i}$  trained on this dataset approximates aleatoric confidence:  $\text{conf}^{\text{alea}}(i, c) \approx f_{\hat{\theta}_i}(\mathbf{A}, \mathbf{X})_{i,c}$ .

**Expected Single Pseudo-Label (ESP).** The right-hand side of Theorem 2.6 is equivalent to epistemic uncertainty and can also be approximated as a proxy for acquisition. For both the numerator and denominator, we estimate the joint probabilities over  $\mathbf{y}_{\mathcal{U}-i}$  as a product of marginal probabilities given by a classifier. The probability of the denominator, can be calculated directly from the predictive distribution of  $f_{\theta}$  (see Appendix J):  $\mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \hat{\mathbf{y}}_{\mathcal{U}-i} \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}}] \propto f_{\theta}(\mathbf{A}, \mathbf{X})_{i, \hat{\mathbf{y}}_i}^{-1}$

Similarly to MP, we compute the probability in the numerator by training a separate classifier after augmenting the training data with a label for  $\mathbf{y}_i$ . Since for each node  $i$ , we condition only on one additional label  $\mathbf{y}_i$  at a time, we do not use the pseudo-label  $\hat{\mathbf{y}}_i$  but instead take an *expectation* over all realizations  $c \in [C]$  with respect to the predictive distribution of  $f_{\theta}$ .

**Results.** We visualize the AL curves of epistemic US using these approaches in Figure 1 and Appendix J. On all datasets, our proposed approximate epistemic uncertainty outperforms other US techniques, often matching the best-performing non-US AL strategy. We observe that ESP approximations perform better in most cases.

**Table 1:** Average AUC ( $\uparrow$ ) of our proposed approximate ground-truth uncertainty versus existing epistemic US. The best estimator is boldfaced and the runner-up is underlined. Our approximate uncertainty consistently outperforms existing epistemic estimators.

	Ensemble	MC-Dropout	Energy	GPN	BGCN	Ours (MP)	Ours (ESP)
CoraML	63.47	59.17	63.97	54.75	44.45	<u>68.16</u>	<b>71.45</b>
Citeseer	82.94	78.86	81.59	65.31	58.68	<b>83.96</b>	83.43
Pubmed	<u>63.70</u>	58.67	59.64	58.82	55.19	n/f	<b>64.36</b>
Photos	<u>84.46</u>	72.42	74.66	54.78	70.83	75.37	<b>85.52</b>
Computers	<u>68.38</u>	51.02	59.62	39.21	58.64	n/f	<b>72.54</b>

While both algorithms estimate the same quantity they rely on different approximations. Therefore, it is expected that they behave differently in practice. While MP substitutes the labels of all unobserved nodes  $\mathcal{U}$  in the dataset with pseudo-labels simultaneously, the ESP method only considers one pseudo-label at a time and factors in the belief of  $f_{\theta}$  by taking an expectation over possible labels. This makes it less prone to erroneous pseudo-labels which, especially for small training sets, are to be expected. Both MP and ESP approximations are also exposed to miscalibration of  $f_{\theta}$  and  $f_{\hat{\theta}}$ , a well-known issue in GNN training [15]. As we discussed in Section 2, message passing GNNs may also not be suited to model the underlying data distribution well. Despite these limitations, approximating and disentangling uncertainty by following our theoretical analysis elevates US from often sub-random to competitive performance levels.

While the performance our approximate approach to ground-truth uncertainty consistently outperforms all existing US methods, our main point is not to propose this as a novel AL strategy. Instead, these results underline the high applicability of our theoretical findings to the development of uncertainty estimators on graphs. We show that epistemic uncertainty is a provably well-suited acquisition proxy for AL on graphs. It identifies highly informative queries in practice when taking into account the theoretical insights of our work. We remark that the paradigm of estimating uncertainty by approximating an unknown generative process may be particularly useful for the domain of graphs as it enables a principled of how interdependence should be accounted for: The non-i.i.d. nature of the problem is implicitly accounted for by estimating a non-i.i.d. data distribution.

## 4 Limitations

Our work highlights the potential of US for AL on graphs and the role of epistemic uncertainty as an optimal guide. While we propose an effective off-the-shelf approximation of our theoretical analysis our goal is not primarily to provide an efficient AL strategy to be deployed in-practice. We leave the development of uncertainty estimators that align with our theoretical work to future work.

## 5 Conclusion

Our study sheds light on the potential and challenges of AL on graphs. We introduce ground-truth uncertainty estimates for node classification and prove the alignment of US with AL. An empirical analysis on real data shows that epistemic US is a highly effective AL strategy when uncertainty is properly disentangled compared to existing strategies. We believe this to be a highly relevant result for uncertainty estimation on graphs which, so far, neglected AL as an evaluation criterion. Our work, thus, lays the necessary theoretical groundwork for developing uncertainty estimators on graphs.

## Impact Statement

Our paper aims at advancing the field of uncertainty estimation on interdependent data which contributes to making AI systems more reliable. By formally establishing a connection to Active Learning, we enable novel acquisition strategies in the domain of graphs to be more data efficient which in turn can help in reducing their negative environmental impact. Nonetheless, advances in the field of Machine Learning are always exposed to non-benign usage. While we acknowledge the various well-established societal and ethical implications of such research we believe the potential positive impacts of our paper outweigh these risks. We encourage both researchers and practitioners building on our analysis to take special care not only for potential misuse but also to thoroughly validate uncertainty estimators that our work proposes, especially in safety-critical domains.

## References

- [1] William H. Beluch, Tim Genewein, Andreas Nurnberger, and Jan M. Kohler. The power of ensembles for active learning in image classification. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9368–9377, 2018. doi: 10.1109/CVPR.2018.00976. 1, 13, 14
- [2] Ajay J. Joshi, Fatih Murat Porikli, and Nikolaos Papanikolopoulos. Multi-class active learning for image classification. *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2372–2379, 2009. 1
- [3] Armen Der Kiureghian and Ove Ditlevsen. Aleatory or epistemic? does it matter? *Structural Safety*, 31:105–112, 2009. ISSN 0167-4730. doi: <https://doi.org/10.1016/j.strusafe.2008.06.020>. 1, 9
- [4] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision?, 2017. 1, 9, 11
- [5] Vu-Linh Nguyen, Mohammad Shaker, and Eyke Hüllermeier. How to measure uncertainty in uncertainty sampling for active learning. *Machine Learning*, 111, 2022. doi: 10.1007/s10994-021-06003-9. 1, 11, 13
- [6] Yarin Gal, Riashat Islam, and Zoubin Ghahramani. Deep bayesian active learning with image data. *CoRR*, abs/1703.02910, 2017. 13, 14
- [7] Andreas Kirsch, Joost van Amersfoort, and Yarin Gal. Batchbald: Efficient and diverse batch acquisition for deep bayesian active learning, 2019. 12, 13, 14
- [8] Sebastian Schmidt and Stephan Günnemann. Stream-based active learning by exploiting temporal properties in perception with temporal predicted loss, 2023. 1, 11
- [9] Xiaojin Zhu, John Lafferty, and Zoubin Ghahramani. Combining active learning and semi-supervised learning using gaussian fields and harmonic functions. In *ICML 2003 workshop on the continuum from labeled to unlabeled data in machine learning and data mining*, volume 3, 2003. 1, 11
- [10] Kwang-Sung Jun and Robert Nowak. Graph-based active learning: A new look at expected error minimization. In *2016 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, pages 1325–1329. IEEE, 2016. 11
- [11] Florence Regol, Soumyasundar Pal, Yingxue Zhang, and Mark Coates. Active learning on attributed graphs via graph cognizant logistic regression and preemptive query generation. In *International Conference on Machine Learning*, pages 8041–8050. PMLR, 2020. 11, 12, 14, 15



- [12] Yuexin Wu, Yichong Xu, Aarti Singh, Yiming Yang, and Artur Dubrawski. Active learning for graph neural networks via node feature propagation, 2021.
- [13] Wentao Zhang, Yexin Wang, Zhenbang You, Meng Cao, Ping Huang, Jiulong Shan, Zhi Yang, and Bin Cui. Information gain propagation: A new way to graph active learning with soft labels. In *ICLR*, 2022. 1
- [14] Maximilian Stadler, Bertrand Charpentier, Simon Geisler, Daniel Zügner, and Stephan Günnemann. Graph posterior network: Bayesian predictive uncertainty for node classification. *Advances in Neural Information Processing Systems*, 34:18033–18048, 2021. 1, 11, 13, 14, 15
- [15] Hans Hao-Hsun Hsu, Yuesong Shen, Christian Tomani, and Daniel Cremers. What makes graph neural networks miscalibrated? *Advances in Neural Information Processing Systems*, 35: 13775–13786, 2022. 4, 23
- [16] Stefan Depeweg, José Miguel Hernández-Lobato, Finale Doshi-Velez, and Steffen Udluft. Decomposition of uncertainty in bayesian deep learning for efficient and risk-sensitive learning, 2018. 9
- [17] Pengzhen Ren, Yun Xiao, Xiaojun Chang, Po-Yao Huang, Zhihui Li, Brij B. Gupta, Xiaojiang Chen, and Xin Wang. A survey of deep active learning, 2021. 11
- [18] Xueying Zhan, Qingzhong Wang, Kuan-hao Huang, Haoyi Xiong, Dejing Dou, and Antoni B. Chan. A Comparative Survey of Deep Active Learning, 2022. 11
- [19] Ozan Sener and Silvio Savarese. Active learning for convolutional neural networks: A core-set approach, 2018. 11, 12
- [20] Samarth Sinha, Sayna Ebrahimi, and Trevor Darrell. Variational Adversarial Active Learning, 2019. 11
- [21] Changjian Shui, Fan Zhou, Christian C Gagné, and Boyu Wang. Deep Active Learning: Unified and Principled Method for Query and Training, 2020. 11
- [22] C. E. Shannon. A mathematical theory of communication, 1948. 11
- [23] Dan Wang and Yi Shang. A new active labeling method for deep learning. In *2014 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2014. doi: 10.1109/IJCNN.2014.6889457. 11
- [24] Neil Houlsby, Ferenc Huszár, Zoubin Ghahramani, and Márton Lengyel. Bayesian Active Learning for Classification and Preference Learning, 2011. 11
- [25] Donggeun Yoo and In So Kweon. Learning loss for active learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 11
- [26] Vu-Linh Nguyen, Bastien Destercke, and Eyke Höllermeier. Epistemic Uncertainty Sampling, 2019. 11
- [27] Changchang Yin, Buyue Qian, Shilei Cao, Xiaoyu Li, Jishang Wei, Qinghua Zheng, and Ian Davidson. Deep Similarity-Based Batch Mode Active Learning with Exploration-Exploitation. In *2017 IEEE International Conference on Data Mining (ICDM)*, 2017. doi: 10.1109/ICDM.2017.67. 11
- [28] Jordan T. Ash, Chicheng Zhang, Akshay Krishnamurthy, John Langford, and Alekh Agarwal. Deep Batch Active Learning by Diverse, Uncertain Gradient Lower Bounds, 2020. 11
- [29] Jordan T Ash, Chicheng Zhang, Akshay Krishnamurthy, John Langford, and Alekh Agarwal. Deep batch active learning by diverse, uncertain gradient lower bounds. *arXiv preprint arXiv:1906.03671*, 2019. 11
- [30] Gautam Dasarathy, Robert Nowak, and Xiaojin Zhu. S2: An efficient graph based active learning algorithm with application to nonparametric classification. In Peter Grünwald, Elad Hazan, and Satyen Kale, editors, *Proceedings of The 28th Conference on Learning Theory*, volume 40 of *Proceedings of Machine Learning Research*, pages 503–522, Paris, France, 03–06 Jul 2015. PMLR. URL <https://proceedings.mlr.press/v40/Dasarathy15.html>. 11
- [31] Jifan Zhang, Julian Katz-Samuels, and Robert Nowak. GALAXY: Graph-based active learning at the extreme. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 26223–26238.

- PMLR, 17–23 Jul 2022. URL <https://proceedings.mlr.press/v162/zhang22k.html>. 11
- [32] Yifei Ma, Roman Garnett, and Jeff Schneider.  $\sigma$ -optimality for active learning on gaussian random fields. *Advances in Neural Information Processing Systems*, 26, 2013. 11
  - [33] Ming Ji and Jiawei Han. A variance minimization criterion to active learning on graphs. In *Artificial Intelligence and Statistics*, pages 556–564. PMLR, 2012.
  - [34] Dimitris Berberidis and Georgios B Giannakis. Data-adaptive active sampling for efficient graph-cognizant classification. *IEEE Transactions on Signal Processing*, 66:5167–5179, 2018. 11
  - [35] Yayong Li, Jie Yin, and Ling Chen. Seal: Semisupervised adversarial active learning on attributed graphs. *IEEE Transactions on Neural Networks and Learning Systems*, 32:3136–3147, 2020. 11, 13, 15
  - [36] Shengding Hu, Zheng Xiong, Meng Qu, Xingdi Yuan, Marc-Alexandre Côté, Zhiyuan Liu, and Jian Tang. Graph policy network for transferable active learning on graphs, 2020. 11
  - [37] Wentao Zhang, Yexin Wang, Zhenbang You, Meng Cao, Ping Huang, Jiulong Shan, Zhi Yang, and Bin Cui. Information gain propagation: a new way to graph active learning with soft labels, 2022. 11
  - [38] Kaushalya Madhawa and Tsuyoshi Murata. Active learning for node classification: An evaluation. *Entropy*, 22, 2020. 11, 12
  - [39] Hongyun Cai, Vincent W. Zheng, and Kevin Chen-Chuan Chang. Active learning for graph embedding, 2017. 11, 12
  - [40] Li Gao, Hong Yang, Chuan Zhou, Jia Wu, Shirui Pan, and Yue Hu. Active discriminative network representation learning. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pages 2142–2148. International Joint Conferences on Artificial Intelligence Organization, 2018. doi: 10.24963/ijcai.2018/296. 12
  - [41] Xiaoting Li, Yuhang Wu, Vineeth Rakesh, Yusan Lin, Hao Yang, and Fei Wang. Smartquery: An active learning framework for graph neural networks through hybrid uncertainty reduction. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, pages 4199–4203, 2022. 11
  - [42] Manali Sharma and Mustafa Bilgic. Evidence-based uncertainty sampling for active learning. *Data Mining and Knowledge Discovery*, 31:164–202, 2016. 11
  - [43] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks, 2017. 11, 14, 15
  - [44] Johannes Klicpera, Aleksandar Bojchevski, and Stephan Günnemann. Personalized embedding propagation: Combining neural networks on graphs with personalized pagerank. *CoRR*, abs/1810.05997, 2018. 11, 15
  - [45] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles, 2017. 11
  - [46] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning, 2016. 11
  - [47] Yu Rong, Wenbing Huang, Tingyang Xu, and Junzhou Huang. Dropedge: Towards deep graph convolutional networks on node classification. In *International Conference on Learning Representations*, 2020. 11
  - [48] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural networks, 2015. 11, 15
  - [49] Jakob Gawlikowski, Cedrique Rovile Njiteucheu Tassi, Mohsin Ali, Jongseok Lee, Matthias Humt, Jianxiang Feng, Anna Kruspe, Rudolph Triebel, Peter Jung, Ribana Roscher, et al. A survey of uncertainty in deep neural networks. *Artificial Intelligence Review*, pages 1–77, 2023. 11
  - [50] Tom Wollschläger, Nicholas Gao, Bertrand Charpentier, Mohamed Amine Ketata, and Stephan Günnemann. Uncertainty estimation for molecules: Desiderata and methods, 2023. 12
  - [51] Zhao-Yang Liu, Shao-Yuan Li, Songcan Chen, Yao Hu, and Sheng-Jun Huang. Uncertainty aware graph gaussian process for semi-supervised learning. 34, 2020. 12

- [52] Viacheslav Borovitskiy, Iskander Azangulov, Alexander Terenin, Peter Mostowsky, Marc Deisenroth, and Nicolas Durrande. Matérn gaussian processes on graphs. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130, pages 2593–2601. PMLR, 2021. [12](#)
- [53] Sanghamitra Bandyopadhyay, Ujjwal Maulik, Lawrence B Holder, Diane J Cook, and Lise Getoor. Link-based classification. *Advanced methods for knowledge discovery from complex data*, pages 189–207, 2005. [12](#), [14](#)
- [54] Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-Rad. Collective classification in network data. *AI magazine*, 29:93–93, 2008. [12](#)
- [55] C Lee Giles, Kurt D Bollacker, and Steve Lawrence. Citeseer: An automatic citation indexing system. In *Proceedings of the third ACM conference on Digital libraries*, pages 89–98, 1998. [12](#), [14](#)
- [56] Galileo Namata, Ben London, Lise Getoor, Bert Huang, and U Edu. Query-driven active surveying for collective classification. In *10th international workshop on mining and learning with graphs*, volume 8, page 1, 2012. [12](#), [14](#)
- [57] Julian McAuley, Christopher Targett, Qinfeng Shi, and Anton Van Den Hengel. Image-based recommendations on styles and substitutes. In *Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval*, pages 43–52, 2015. [12](#)
- [58] Felix Wu, Amauri Souza, Tianyi Zhang, Christopher Fifty, Tao Yu, and Kilian Weinberger. Simplifying graph convolutional networks. In *International conference on machine learning*, pages 6861–6871. PMLR, 2019. [13](#), [24](#)
- [59] Weitang Liu, Xiaoyun Wang, John D. Owens, and Yixuan Li. Energy-based out-of-distribution detection, 2021. [13](#)
- [60] Qitian Wu, Yiting Chen, Chenxiao Yang, and Junchi Yan. Energy-based out-of-distribution detection for graph neural networks, 2023. [13](#), [14](#)
- [61] Andrew Kachites McCallum, Kamal Nigam, Jason Rennie, and Kristie Seymore. Automating the construction of internet portals with machine learning. *Information Retrieval*, 3:127–163, 2000. [14](#)
- [62] Yash Deshpande, Andrea Montanari, Elchanan Mossel, and Subhabrata Sen. Contextual stochastic block models. In *Neural Information Processing Systems*, 2018. [14](#)
- [63] Lukas Gosch, Daniel Sturm, Simon Geisler, and Stephan Günnemann. Revisiting robustness in graph machine learning, 2023. [15](#)
- [64] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [15](#)
- [65] Mahendra Mariadassou, Stéphane Robin, and Corinne Vacher. Uncovering latent structure in valued graphs: a variational approach. 2010. [17](#)
- [66] Tommi S Jaakkola. Tutorial on variational approximation methods. *Advanced mean field methods: theory and practice*, pages 129–159, 2000. [17](#), [18](#)
- [67] John Palowitch, Anton Tsitsulin, Brandon Mayer, and Bryan Perozzi. GraphWorld. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. ACM, 2022. doi: 10.1145/3534678.3539203. [18](#)
- [68] Will Grathwohl, Kuan-Chieh Wang, Jörn-Henrik Jacobsen, David Duvenaud, Mohammad Norouzi, and Kevin Swersky. Your classifier is secretly an energy based model and you should treat it like one. *arXiv preprint arXiv:1912.03263*, 2019. [24](#)



## A Background

**AL for Semi-Supervised Node classification.** Let  $G = (\mathcal{V}, \mathcal{E})$ , be a graph with a set of nodes  $\mathcal{V}$  and edges  $\mathcal{E} \subseteq \{\{i, j\} : i, j \in \mathcal{V}\}$ . With  $n = |\mathcal{V}|$  number of nodes the adjacency matrix is defined as  $\mathbf{A} \in \{0, 1\}^{n \times n}$ , where  $A_{ij} = A_{ji} = 1$  if there exists an edge between node  $i$  and node  $j$ . The features are represented as matrix  $\mathbf{X} \in \mathbb{R}^{n \times d}$ . In node classification, every node has a label represented by a vector  $\mathbf{y} \in [C]^n$ . We decompose the set of all nodes into  $\mathcal{V} = \mathcal{U} \cup \mathcal{O}$  where  $\mathcal{O}$  contains nodes with observed labels  $\mathbf{y}_{\mathcal{O}}$  and  $\mathcal{U}$  contains remaining unobserved nodes  $\mathbf{y}_{\mathcal{U}}$  that we want to infer, where  $\mathcal{V} = \mathcal{U} \cup \mathcal{O}$ . We consider pool-based AL: In each iteration, the learner queries an oracle for a label  $\mathbf{y}_i$  from the set of unobserved labels  $\mathbf{y}_{\mathcal{U}}$ , adds it to the set of observed labels  $\mathbf{y}_{\mathcal{O}}$  and retrains the model.

**Uncertainty in Machine Learning.** Alongside the predicted label  $\mathbf{y}_i$ , it is crucial to consider the associated uncertainty [3, 4]. Aleatoric uncertainty  $u^{\text{alea}}$  is the inherent uncertainty that comes from elements like experimental randomness. It can not be reduced by acquiring more data. On the other hand, Epistemic uncertainty  $u^{\text{epi}}$  reflects knowledge gaps, which can be addressed by data acquisition. A commonly employed conceptualization [16] defines the total predictive uncertainty  $u^{\text{total}}$  as encapsulating both reducible and irreducible factors, i.e.  $u^{\text{total}} = u^{\text{epi}} + u^{\text{alea}}$ .

## B Proofs

*Theorem 1.* Epistemic uncertainty  $u^{\text{epi}}(i, \mathbf{y}_i^{\text{gt}})$  of a node  $i$  is equivalent to the relative gain its acquisition provides to the posterior over the remaining true labels:

$$u^{\text{epi}}(i, \mathbf{y}_i^{\text{gt}}) = \frac{\mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}}, \mathbf{y}_i = \mathbf{y}_i^{\text{gt}}]}{\mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}}]}$$

Hence, acquiring the most epistemically uncertain node is an optimal AL strategy for  $f_{\theta}^*$ .

*Proof.*

$$\begin{aligned} \text{conf}^{\text{alea}}(i, \mathbf{y}_i^{\text{gt}}) &= \int \mathbb{P}[\mathbf{y}_i = \mathbf{y}_i^{\text{gt}} \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{-i} = \mathbf{y}_{-i}^{\text{gt}}, \theta] p(\theta \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{-i} = \mathbf{y}_{-i}^{\text{gt}}) d\theta \\ &= \mathbb{P}[\mathbf{y}_i = \mathbf{y}_i^{\text{gt}} \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{-i} = \mathbf{y}_{-i}^{\text{gt}}] \\ &= \mathbb{P}[\mathbf{y}_i = \mathbf{y}_i^{\text{gt}} \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}}, \mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}}] \\ &= \frac{\mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}}, \mathbf{y}_i = \mathbf{y}_i^{\text{gt}}]}{\mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}}]} \mathbb{P}[\mathbf{y}_i = \mathbf{y}_i^{\text{gt}} \mid \mathbf{X}, \mathbf{A}, \mathbf{y}_{\mathcal{O}}] \\ &= \frac{\mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}}, \mathbf{y}_i = \mathbf{y}_i^{\text{gt}}]}{\mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}}]} \text{conf}^{\text{total}}(i, \mathbf{y}_i^{\text{gt}}) \\ u^{\text{epi}}(i, \mathbf{y}_i^{\text{gt}}) &= \frac{\text{conf}^{\text{alea}}(i, \mathbf{y}_i^{\text{gt}})}{\text{conf}^{\text{total}}(i, \mathbf{y}_i^{\text{gt}})} = \frac{\mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}}, \mathbf{y}_i = \mathbf{y}_i^{\text{gt}}]}{\mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}}]} \end{aligned}$$

First, we insert our definition of aleatoric confidence and marginalize  $\theta$ . Then, we split  $\mathbf{y}_{-i}$  into two parts: observed  $\mathbf{y}_{\mathcal{O}}$  and unobserved  $\mathbf{y}_{\mathcal{U}-i}$ . As we assign the ground-truth labels to both, the exact partition is not relevant. Next, we use Bayes law to get a distribution over the unobserved node labels  $\mathbf{y}_{\mathcal{U}}$ . Similarly to the first step, we marginalize  $\theta$  to obtain  $\text{conf}^{\text{total}}(i, \mathbf{y}_i^{\text{gt}})$ . In the last step, we see that the right term matches our definition of total confidence  $\text{conf}^{\text{total}}$ .  $\square$

*Proposition 1.* Total uncertainty  $u^{\text{total}}(i, \mathbf{y}_i^{\text{gt}})$  of a node  $i$  is proportional to the posterior over the unobserved true labels  $\mathbf{y}_{\mathcal{U}-i}^{\text{gt}}$  after acquiring its label  $\mathbf{y}_i^{\text{gt}}$ :

$$u^{\text{total}}(i, \mathbf{y}_i^{\text{gt}}) \propto \mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}}, \mathbf{y}_i = \mathbf{y}_i^{\text{gt}}]$$

*Proof.*

$$\begin{aligned}
 u^{\text{total}}(i, \mathbf{y}_i^{\text{gt}}) &= \frac{1}{\text{conf}^{\text{total}}(i, \mathbf{y}_i^{\text{gt}})} \\
 &= \frac{\text{conf}^{\text{alea}}(i, \mathbf{y}_i^{\text{gt}})}{\text{conf}^{\text{total}}(i, \mathbf{y}_i^{\text{gt}})} \frac{1}{\text{conf}^{\text{alea}}(i, \mathbf{y}_i^{\text{gt}})} \\
 &= \frac{\mathbb{P}[\mathbf{y}_{\mathcal{U}} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \dots, \mathbf{y}_i = \mathbf{y}_i^{\text{gt}}]}{\mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \dots]} \frac{1}{\text{conf}^{\text{alea}}(i, \mathbf{y}_i^{\text{gt}})} \\
 &= \frac{\mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \dots, \mathbf{y}_i = \mathbf{y}_i^{\text{gt}}]}{\mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \dots]} \frac{1}{\mathbb{P}[\mathbf{y}_i = \mathbf{y}_i^{\text{gt}} \mid \mathbf{y}_{\mathcal{U}} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}}, \dots]} \\
 &= \frac{\mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \dots, \mathbf{y}_i = \mathbf{y}_i^{\text{gt}}]}{\mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \dots]} \frac{\mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \dots]}{\mathbb{P}[\mathbf{y}_i = \mathbf{y}_i^{\text{gt}}, \mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \dots]} \\
 &= \frac{\mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \dots, \mathbf{y}_i = \mathbf{y}_i^{\text{gt}}]}{\mathbb{P}[\mathbf{y}_i = \mathbf{y}_i^{\text{gt}} \mid \mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \dots]} \\
 &= \frac{\mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \dots, \mathbf{y}_i = \mathbf{y}_i^{\text{gt}}]}{\mathbb{P}[\mathbf{y}_{\mathcal{U}} = \mathbf{y}_{\mathcal{U}}^{\text{gt}} \mid \dots]} \\
 &\propto \mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \dots, \mathbf{y}_i = \mathbf{y}_i^{\text{gt}}]
 \end{aligned}$$

Here, we abbreviated  $\mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}}$  with  $\dots$  for clarity. First, we insert the result from Theorem 2.6. Then, we use the definition of conditional probability to replace the inverse of the aleatoric uncertainty. After canceling terms, we observe that  $\mathbb{P}[\mathbf{y}_{\mathcal{U}} = \mathbf{y}_{\mathcal{U}}^{\text{gt}} \mid \dots]$  is a constant with respect to  $i$  and arrive at the desired result that the total uncertainty is proportional to the posterior over the true labels  $\mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \dots, \mathbf{y}_i = \mathbf{y}_i^{\text{gt}}]$ . As in Proposition H.1, we implicitly marginalized the learnable parameters  $\theta$  of the Bayesian classifier.  $\square$

*Proposition 2.* Aleatoric uncertainty  $u^{\text{alea}}(i, \mathbf{y}_i^{\text{gt}})$  of a node  $i$  is proportional to the posterior over the unobserved true labels  $\mathbf{y}_{\mathcal{U}-i}^{\text{gt}}$  without acquiring its label  $\mathbf{y}_i^{\text{gt}}$ :

$$u^{\text{alea}}(i, \mathbf{y}_i^{\text{gt}}) \propto \mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}}]$$

*Proof.*

$$\begin{aligned}
 u^{\text{alea}}(i, \mathbf{y}_i^{\text{gt}}) &= \frac{1}{\text{conf}^{\text{alea}}(i, \mathbf{y}_i^{\text{gt}})} \\
 &= \frac{\text{conf}^{\text{total}}(i, \mathbf{y}_i^{\text{gt}})}{\text{conf}^{\text{alea}}(i, \mathbf{y}_i^{\text{gt}})} \frac{1}{\text{conf}^{\text{total}}(i, \mathbf{y}_i^{\text{gt}})} \\
 &= \frac{\mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \dots]}{\mathbb{P}[\mathbf{y}_{\mathcal{U}} = \mathbf{y}_{\mathcal{U}}^{\text{gt}} \mid \dots, \mathbf{y}_i = \mathbf{y}_i^{\text{gt}}]} \frac{1}{\text{conf}^{\text{total}}(i, \mathbf{y}_i^{\text{gt}})} \\
 &= \frac{\mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \dots]}{\mathbb{P}[\mathbf{y}_{\mathcal{U}} = \mathbf{y}_{\mathcal{U}}^{\text{gt}} \mid \dots, \mathbf{y}_i = \mathbf{y}_i^{\text{gt}}]} \frac{1}{\mathbb{P}[\mathbf{y}_i = \mathbf{y}_i^{\text{gt}} \mid \dots]} \\
 &= \frac{\mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \dots]}{\mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}}, \mathbf{y}_i = \mathbf{y}_i^{\text{gt}} \mid \dots]} \\
 &= \frac{\mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \dots]}{\mathbb{P}[\mathbf{y}_{\mathcal{U}} = \mathbf{y}_{\mathcal{U}}^{\text{gt}} \mid \dots]} \\
 &\propto \mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \dots]
 \end{aligned}$$

Again, we abbreviated  $\mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}}$  with  $\dots$  for clarity. First, we insert the result from Theorem 2.6 into  $\text{conf}^{\text{total}}(i, \mathbf{y}_i^{\text{gt}})/\text{conf}^{\text{alea}}(i, \mathbf{y}_i^{\text{gt}})$ . Then, insert the definition of total confidence. We apply the

law of conditional probability, we see that  $\mathbb{P}[\mathbf{y}_{\mathcal{U}} = \mathbf{y}_{\mathcal{U}}^{\text{gt}} \mid \dots]$  is a constant with respect to  $i$  and arrive at the desired result that the aleatoric uncertainty is proportional to the posterior over the true labels  $\mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \dots]$ .


□

## C Related Work

**Active Learning on Independent and Identically Distributed Data.** AL has seen substantial exploration in the context of i.i.d. data [17]. Approaches can be divided into three main categories: diversity-based, uncertainty-based, or a combination thereof [18]. Diversity or representation-based methods query data samples that best represent the full dataset, i.e. they opt for a diverse set of data points. Approaches like KMeans or Coreset opt to minimize the difference in model loss between the selected training set and the whole data set [19]. Other approaches use adversarial techniques to estimate the representativeness and diversity of new samples [20, 21]. Uncertainty-based approaches query instances that the classifier is most uncertain about. It is commonly computed from the predictive distribution of a classifier calculating its entropy [22], the margin between the two most likely labels, or as the probability of the least confident label [23]. Houlsby et al. [24] introduce Bayesian Active Learning by Disagreement (BALD), which queries points with high mutual information between the model parameters and the class label. Another line of work uses a loss-prediction module as a proxy for uncertainty [8, 25]. Nguyen et al. [26] leverage disentangled uncertainties [4] for AL and propose that epistemic uncertainty is a better proxy for US than aleatoric estimates. Other approaches linearly combine diversity- and uncertainty-based measures [27] or employ a two-step optimization scheme [18, 28]. BADGE [29] queries diverse and uncertain instances in the gradient space of the model using clustering. Another line of work constructs similarity graphs between instances to enforce diversity among queries [30]. GALAXY [31] proposes further refines this approach by mitigating class imbalance. In contrast, our work concerns settings where the adjacency matrix  $A$  is explicitly given.

**Active Learning on Interdependent Graph Data.** While a plethora of studies exist on AL for i.i.d. data, only a limited amount of work addresses interdependent data like graphs. Previous methods approach AL on graphs from different perspectives, including random fields [9, 32–34], risk minimization [10, 11], adversarial learning [35], knowledge transfer between graphs [36] or querying cheap soft labels [37]. US is typically only considered in terms of the predictive distribution [38] which does not disentangle aleatoric and epistemic components. Also, other uncertainty-reliant approaches do not make that distinction [39–41] even though the literature on i.i.d. data suggests that US benefits from disentangled uncertainty estimators [5, 42]. The exploration of US for AL on graphs beyond total uncertainty remains uncharted territory. Our work targets this gap and showcases the unrealized potential of epistemic uncertainty estimators for AL on graph data.

**Uncertainty Estimation on Graphs.** US strategies necessitate accurate uncertainty estimates which can be obtained in different ways. In classification, the predictive distribution of deterministic classifiers has been used to obtain aleatoric uncertainty [14]. The Graph Convolutional Network (GCN) [43] iteratively updates the representation  $H^{(l)}$  of each node using a linear transformation  $W^{(l)}$  and subsequent diffusion along the edges and non-linearity  $\sigma$ . Formally, a GCN layer is expressed as  $f(H^{(l)}, A) = \sigma(AH^{(l)}W^{(l)})$ . The APPNP model [44] first transforms the nodes features independently and then diffuses predictions according to approximate Personalized Page Rank (PPR) scores. Bayesian approaches model the posterior distribution over the model parameters. **Ensembles** [45] fall into this category. They approximate the posterior over model parameters through a collection of independently trained models. **Monte Carlo Dropout** (MC-Dropout) [46] instead emulates a distribution over model parameters by applying dropout at inference time. DropEdge [47] proposed to additionally drop edges to reduce over-fitting and over-smoothing. Variational Bayes methods place a prior on the model parameters (BGCN) and sample different parameter sets for each forward pass [48]. They allow access to disentangled uncertainty estimates by approximating a distribution over predictions from multiple forward passes. Commonly employed measures of epistemic uncertainty are the mutual information between the model weights and predicted labels [49] or the variance in confidence about the predicted label [14]. Evidential methods like GPN [14] disentangle epistemic and aleatoric uncertainty by outputting the parameters of a Dirichlet prior to the categorical predictive distribution. This method has shown strong performance in detecting distribution shifts. Finally, Gaussian processes on graphs, while potentially leading to



**Figure 2:** Accuracy of AL strategies on Citeseer using a GCN (left) / SGC (right) classifier. Except for GEEM, which is only tractable for SGCs, traditional AL can not significantly outperform random selection.

strong uncertainty estimates in specific domains [50], they do not disentangle aleatoric and epistemic uncertainty [51, 52].

## D Benchmarking Uncertainty Sampling Approaches for Active Learning on Graphs

Previous studies on non-uncertainty-based AL on graphs find that AL strategies struggle to consistently outperform random sampling [38]. We, therefore, ask the question of whether US shows any merit when using state-of-the-art uncertainty estimation. We are the first to design a comprehensive AL benchmark for node classification that not only considers traditional methods but also includes a variety of uncertainty estimators for US.


**Experimental Setup.** We evaluate AL on five common citation benchmark datasets for node classification: **CoramL** [53], **Citeseer** [54, 55], **PubMed** [56] as well as the co-purchase graphs **Amazon Photos** and **Amazon Computers** [57]. We evaluate the models of Appendix C: **GCN**, **APPNP**, **MC-Dropout**, **BGCN**, **GPN** and **Ensembles** and report average results over multiple AL runs (see Appendix E).

When acquiring multiple labels per iteration, greedily selecting the most promising candidates might overestimate the performance improvement [7]. We therefore only acquire a single label in each iteration, enabling us to analyze the performance of different acquisition strategies without having to consider the potential side effects of batched acquisition. We initially label one node per class and fix the acquisition budget to  $4C$ .

In addition to a qualitative evaluation of the accuracy curves of different strategies in Figure 2, we report the accuracy after the labeling budget is exhausted in Table 6. Good acquisition functions should achieve higher accuracy at a lower amount of queries, which in turn results in a larger area under (AUC) the visualized curves. After normalization, this metric quantifies the average accuracy which we report as a summary of the AL performance in Table 2.

We proceed as follows: First, we benchmark traditional AL approaches and find that only one consistently outperforms **Random** selection that acquires labels with uniform probability. Then, we show that US fails to even match the performance of random queries in many instances, contrasting its successful application to i.i.d. data. We supply the corresponding AUC and final accuracy scores as well as visualizations on all datasets in Appendix F.

**Non-Uncertainty-based Strategies.** We first examine strategies that do not exclusively rely on uncertainty. (i) **Coreset** [19] opts to find a core-set cover of the training pool by selecting nodes that maximize the minimal distance in the latent space of a classifier to already acquired instances. (ii) **Coreset-PPR** is similar to Coreset, but we use inverse Personalized Page Rank (PPR) scores as a distance measure to select structurally dissimilar nodes. (iii) **Coreset Features**. Distances between nodes are computed only in terms of input features. (iv) **Degree** and **PPR**. We acquire nodes with the highest corresponding centrality measure. (v) **AGE** [39] and **ANRMAB** [40] combine total predictive uncertainty, informativeness, and representativeness metrics. (vi) **GEEM** [11] uses risk minimization to select the next query. Because of its high computational cost, we follow its authors and employ an



**Figure 3:** US on Citeseer. No method significantly outperforms random selection.

SGC [58] backbone. (vii) **SEAL** [35] uses adversarial learning to identify nodes dissimilar from the labeled set. If applicable, we also consider setting  $\mathcal{A} = \mathcal{I}$  to exclude structural information similar to [14].

**Observations:** Figure 2 shows that only GEEM identifies a training set that is significantly more informative than random sampling. The performance of GEEM, however, comes at a high computational cost, as it requires training  $\mathcal{O}(nC)$  models in each acquisition which makes it intractable for larger datasets and models beyond SGC. Structural Coreset strategies work well on both co-purchase networks but do not show strong results on other graphs. This highlights a notable gap between many commonly used acquisition functions and a potential optimal strategy.

**Uncertainty Sampling.** We evaluate US using different uncertainty estimators: (i) **Aleatoric.** Like [14], we compute  $u^{\text{alea}} = -\max_c p_c$ . (ii) **Epistemic.** For models perform multiple predictions (MC-Dropout, Ensembles, BGCNs), we compute epistemic uncertainty as the variance of the confidence  $u^{\text{epi}} = \text{Var}[p_{\hat{c}}]$  in the predicted class  $\hat{c}$ . For GPN, we follow the authors and use evidence as a measure of epistemic confidence. (iii) **Energy.** Energy-based models (EBMs) [59, 60] relate uncertainty to the energy  $u = -\tau \log \sum_c \exp(l_i)$  of the predicted logits  $l$ . We apply this estimator to the deterministic GCN, APPNP, and SGC models as a surrogate for epistemic uncertainty. Again, we ablate each strategy by withholding structural information ( $\mathcal{A} = \mathcal{I}$ ).

**Table 2:** Average AUC ( $\uparrow$ ) for different acquisition strategies on different models and datasets. We mark the best strategy per model in bold and underline the runner-up. For each dataset, we highlight the overall best strategy with a  $\dagger$  symbol.

		Baselines	Non-Uncertainty						Uncertainty					
	Inputs		A & X						A	X	A & X		X	
	Model	Random	Coreset	AGE	ANRMAB	GEEM	SEAL	Coreset PPR	Coreset Inputs	Epi/(Energy)	Alea.	Epi/(Energy)	Alea.	
CoraML	GCN	62.51	64.35	64.12	64.24	n/a	<b>66.07</b>	59.53	61.26	63.97	61.30	65.65	64.33	
	APPNP	67.72	67.32	66.12	69.49	n/a	n/a	<b>71.04</b>	64.49	64.92	67.68	69.59	66.69	
	Ensemble	63.89	60.55	64.80	<u>65.10</u>	n/a	n/a	62.65	65.07	63.47	64.03	64.80	<b>65.82</b>	
	MC-Dropout	<b>64.94</b>	64.37	<u>64.44</u>	64.06	n/a	n/a	62.92	64.35	59.17	63.69	61.82	63.87	
	BGCN	45.76	<u>49.37</u>	<b>51.25</b>	47.23	n/a	n/a	39.43	44.85	44.45	46.61	42.74	48.11	
	GPN	56.50	n/a	n/a	n/a	n/a	n/a	<b>58.04</b>	54.02	54.75	57.16	55.89	<u>57.21</u>	
	SGC	63.85	65.23	<u>67.56</u>	61.14	<b>71.39<sup>†</sup></b>	n/a	60.24	59.18	67.51	65.66	65.05	67.13	
Pubmed	GCN	61.56	<u>62.61</u>	<b>69.48</b>	60.31	n/a	58.62	61.71	56.71	59.64	61.85	59.66	60.34	
	APPNP	<u>64.61</u>	63.88	<b>70.18<sup>†</sup></b>	63.83	n/a	n/a	64.21	56.87	63.09	62.37	62.23	63.95	
	Ensemble	59.26	<u>64.25</u>	<b>68.26</b>	60.40	n/a	n/a	61.89	56.36	63.70	61.37	59.71	61.15	
	MC-Dropout	58.30	<u>62.97</u>	<b>65.24</b>	60.50	n/a	n/a	61.43	56.01	58.67	59.07	59.23	62.22	
	BGCN	53.59	<b>59.29</b>	56.93	52.68	n/a	n/a	53.40	51.40	55.19	52.81	<u>57.09</u>	54.62	
	GPN	<u>59.76</u>	n/a	n/a	n/a	n/a	n/a	<b>62.08</b>	54.34	58.82	57.24	56.25	59.37	
	SGC	56.79	64.48	<b>69.20</b>	60.49	<u>64.82</u>	n/a	62.15	52.25	62.04	61.55	61.04	60.74	
AmazonPhotos	GCN	79.06	78.58	75.17	<u>79.97</u>	n/a	71.16	70.20	<b>82.71</b>	74.66	74.61	79.96	79.63	
	APPNP	79.29	<u>81.04</u>	79.02	80.35	n/a	n/a	76.37	<b>84.24</b>	79.72	77.45	80.48	77.69	
	Ensemble	82.23	80.44	77.45	82.77	n/a	n/a	74.93	<u>84.04</u>	<b>84.46</b>	77.85	80.50	81.25	
	MC-Dropout	<u>80.32</u>	76.63	74.75	80.21	n/a	n/a	75.32	<b>82.45</b>	72.42	73.16	69.68	78.80	
	BGCN	71.22	67.15	65.69	70.69	n/a	n/a	59.34	<b>73.39</b>	70.83	67.83	<u>72.21</u>	69.19	
	GPN	62.80	n/a	n/a	n/a	n/a	n/a	55.59	<b>65.07</b>	54.78	60.53	<u>62.90</u>	62.41	
	SGC	80.52	82.32	74.01	80.92	<b>86.43<sup>†</sup></b>	n/a	66.94	<u>84.24</u>	84.01	71.43	80.75	76.38	

**Observations:** In contrast to literature on i.i.d. data [1, 5, 6], we observe none of the US approaches to be effective in Figure 3. While sampling instances with high aleatoric uncertainty matches the performance of random queries, we find epistemic uncertainty to even underperform in many instances. This is surprising, as the efficacy of epistemic US has been demonstrated for i.i.d. data [5, 7]. Only ensemble models match the performance of random sampling and slightly outperform on some datasets. GPN and energy-based approaches can not guide US toward effective queries.



This is an intriguing result as both uncertainty estimators have shown to be highly effective for out-of-distribution detection [14, 60].

## E Experimental Setup

### E.1 Active Learning

We discuss AL on graphs, where given an initial set of labeled instances  $\mathcal{O} \subset \mathcal{V}$  we aim to acquire a set of the unlabeled nodes  $\mathcal{U} \subset \mathcal{V}$  that is optimal in terms of improving the performance of the classifier on the entire graph. That is, we (1) initially label one node randomly drawn from each class. (2) We then re-initialize the model weights and train the classifier until convergence. (3) We employ an acquisition strategy to select one or more unlabeled node(s). (4) We add the acquired label(s) to the training set and repeat the procedure at step (2) until some acquisition budget is exhausted. In contrast to domains where model training is expensive [1, 6, 7], we re-train the classifier after each acquisition iteration. If not stated otherwise, we only acquire one node label in each iteration and fix the acquisition budget to  $5C$ . The resulting final training pools, therefore, contain fewer instances compared to dataset splits commonly used in other work [14, 43].

### E.2 Hyperparameters

As hyperparameter tuning may be unrealistic in AL [11], we do not finetune them on a validation set. One potential strategy to realize hyperparameters is to randomly sample them from the search space over which hyperparameter optimization would be performed. This, however, is expected to lead to notably worse performance for most GNN architectures [11]. Since the goal of our work is to showcase that even under optimal circumstances, contemporary uncertainty estimators can not enable US to be effective, we select one set of hyperparameters reported to be effective by the literature and employ it on all datasets. Specifically, we chose the following values:

Model	Hidden Dimensions	Learning Rate	Max Epochs	Weight Decay	Teleport Probability	Power Iteration Steps	Dropout	Flow Dimension	Number Radial Flow Layers
GCN	[64]	0.001	10,000	0.001	n/a	n/a.	0.0	n/a	n/a.
APPNP	[64]	0.001	10,000	0.001	0.2	10	0.0	n/a	n/a
GPN	[64]	0.001	10,000	0.001	0.2	10	0.5	16	10

**Table 3:** Hyperparameters of different GNN backbones

### E.3 Datasets

**Real-World Datasets.** We evaluate AL approaches on three common node classification benchmark citation datasets: CoraML [53–55, 61], Citeseer [53–55], PubMed [56]. In all datasets, nodes are papers and edges model citations. We consider undirected edges only and select the largest connected component in the dataset. All datasets use bag-of-words representations and we normalize the input features  $\mathbf{x}_i$  node-wise to have a  $L_2$ -norm of 1. We report the statistics of each dataset in Table 4.

Dataset	#Nodes $n$	#Edges $m$	#Input Features $d$	#Classes $c$	Edge Density $\frac{m}{n^2}$	Homophily $p$	Intra-Class Edge Density $p$	Inter-Class Edge Density $q$	SNR $\frac{p}{q}$
CoraML	2,810	15,962	2,879	7	0.20%	78.44%	1.69%	0.06%	28.51
Citeseer	1,681	5,804	602	6	0.20%	92.76%	4.89%	0.03%	141.36
PubMed	19,717	88,648	500	3	0.02%	80.24%	0.07%	0.01%	9.48
AmazonComputers	13,381	491,556	767	10	0.27%	77.72%	3.60%	0.07%	54.74%
AmazonPhotos	7,484	238,086	745	8	42.47%	82.72%	3.77%	0.10%	36.85%

**Table 4:** Dataset statistics.

We additionally compute the edge density  $\frac{m}{n^2}$  as well as the homophily which is the fraction of edges that link between nodes of the same class. For comparability with CSBMs, we also report the average empirical inter-class edge probabilities  $p$  and intra-class edge probabilities  $q$  as well as their ratio, the structural signal-to-noise rate (SNR).

#### Contextual Stochastic Blockmodels (CSBMs).

CSBMs [62] define the following generative process: First, node labels are sampled independently from a prior  $p(\mathbf{y})$ . For each node, features are then drawn independently from a class-

conditional normal distribution  $p(\mathbf{X}_i \mid \mathbf{y}_i) \sim \mathcal{N}(\mu_{\mathbf{y}_i}, \sigma_x^2 \mathbf{I})$ . Each edge in the graph is generated independently according to an affiliation matrix  $\mathbf{F} \in \mathbb{R}^{c \times c}$ , i.e.  $p(\mathbf{A}_{i,j} \mid \mathbf{y}_i, \mathbf{y}_j) \sim \text{Ber}(\mathbf{F}_{\mathbf{y}_i, \mathbf{y}_j})$ . This gives rise to an explicit joint distribution over the graph that factorizes as  $p(\mathbf{A}, \mathbf{X}, \mathbf{y}) = \prod_{i < j} p(\mathbf{A}_{i,j} \mid \mathbf{y}_i, \mathbf{y}_j) \prod_i p(\mathbf{X}_i \mid \mathbf{y}_i) \prod_i p(\mathbf{y}_i)$ .

We generate CSBM graphs with a fixed number of nodes, classes, and input features according to Appendix A. That is, we first sample class labels  $\mathbf{y}_i$  for each node independently from a uniform prior  $\mathbb{P}[\mathbf{y}_i] = 1/C$ . We then create links according to the affiliation matrix  $\mathbf{F}$ , where  $\mathbb{P}[\mathbf{A}_{i,j} = 1 \mid \mathbf{y}_i, \mathbf{y}_j] = \mathbf{F}_{\mathbf{y}_i, \mathbf{y}_j}$ . If not stated otherwise, we use a symmetric and homogeneous affiliation matrix  $\mathbf{F}$ , where we set all diagonal elements to a given intra-class edge probability  $p$  and all off-diagonal elements to a given inter-class edge probability  $q$ . We enforce a structural signal-to-noise ratio (SNR)  $\sigma_A = p/q$  by specifying an expected node degree  $\mathbb{E}[\deg(v)]$  and then solving for:

$$q = \frac{\mathbb{E}[\deg(v)] c}{n - 1} \frac{1}{\sigma_A + c - 1} \quad (5)$$

For each class, we first deterministically draw  $C$  vectors of dimension  $d$  that all have a pairwise distance of  $\delta_X$ . We then use a random rotation to obtain class centers  $\mu_c$ . For each node, we then independently sample its features  $\mathbf{x}_i$  from a normal distribution  $\mathcal{N}(\mathbf{X}_i \mid \mu_{\mathbf{y}_i}, \sigma_X^2 \mathbf{I})$ . We refer to the quotient  $\frac{\delta_X}{\sigma_X}$  as the feature signal-to-noise ratio. The dimension of the feature space is given by  $d = \max(c, \lceil n / \log^2 n \rceil)$ , following [63].

#### E.4 Model Details

If applicable, we use the same GCN [43] backbone for all models. That is, we employ one hidden layer of dimension of 64. The APPNP [44] model uses an MLP with one hidden layer of the same dimension. We diffuse the predictions using 10 steps of power iteration and a teleport probability of 0.2.

If not stated otherwise, ensembles are composed of 10 architecturally identical GCNs. In the case of MCD, we apply dropout with probability 0.1 to each neuron of the GCN backbone independently. The BGCN mimics the GCN but models weights and biases as normal distributions. We follow [48] and regularize these towards a standard normal distribution. Consequentially, we apply a weight of  $\lambda = 0.1$  to the regularization loss.

We follow [14] to configure the GPN model: We use 10 radial flow layers to implement the class-conditional density model with a dimension of 16. For diffusion, we implement the same configuration as for APPNP. Other hyperparameters are set as in [14].

For SEAL, we follow the authors and do not re-train the model after each acquisition [35]. Furthermore, we pick the number of training iterations for the discriminator to be  $n_d = 5$  and the number of iterations for the generator  $n_g = 10$ , a combination that we observed to be successful on one dataset and kept fixed for all others as the authors do not provide them directly and hyperparameter optimization is unrealistic in AL [11].

#### E.5 Training Details

Apart from the GPN, all of the aforementioned models are trained towards the binary-cross-entropy objective using the ADAM [64] optimizer with a learning rate of  $10^{-3}$  and weight decay of  $10^{-3}$ . We also perform early stopping on the validation loss with a patience of 100 iterations. We implement our models in PyTorch and PyTorch Geometric and train on two types of machines: (i) Xeon E5-2630 v4 CPU @ 2.20GHz with a NVIDIA GTX 1080TI GPU and 128 GB of RAM. (ii) AMD EPYC 7543 CPU @ 2.80GHz with a NVIDIA A100 GPU and 128 GB of RAM.

For each dataset, classifier, and acquisition function, we report results averaged over five different dataset splits and five distinct model initializations. In each dataset split, we a priori fix 20% of all nodes as a test set that is reused in every subsequent dataset split and can never be acquired by any strategy.

## F Additional Metrics and Plots

We supply an evaluation of contemporary non-uncertainty-based acquisition strategies on different models as well as various uncertainty estimators for US over all datasets listed in Appendix E.3. We briefly summarize the key insights:

(i) **CoraML**. Figures 4a to 4d show that apart from GEEM no AL strategy is significantly more effective than random sampling. For an SGC classifier, GEEM can identify high-quality training sets. In terms of US approaches, only ensemble methods perform somewhat better than random acquisition. (ii) **Citeseer**. Figures 3, 4e, 4g and 4h show that no AL strategy can match the performance of the GEEM. Again, risk minimization is the strongest performing approach. Only ensembles and energy-based models can compete with random acquisition when concerning epistemic uncertainty estimates. We verify the intriguing trend observed in Appendix D that epistemic uncertainty proxies seem to significantly underperform random queries when being employed for US. This supports our conjecture that contemporary estimators may not properly disentangle or only partially model uncertainty. (iii) **PubMed**. In Figures 4i to 4l, GEEM is outperformed by the AGE baseline. For this dataset, acquiring central nodes (PPR, AGE) appears to be a successful strategy. Interestingly, ANRMAB fails to exploit centrality despite it, in theory, being capable of doing so. US appears to be effective only when employing ensemble methods. (iv) **AmazonPhotos**. On this dataset, GEEM outperforms a centrality-based (PPR) Coreset approach (see Figures 4m to 4p. Concerning US, we again find that only ensemble methods outperform random sampling while other estimators lead to significantly worse performance. (v) **Amazon Computers**. Similar to the other co-purchase network, we observe Coreset PPR to be a strong proxy for AL on the AmazonComputers dataset (Figures 4q to 4t). For this dataset, all US approaches—including ensembles—perform significantly worse than a random strategy.

This affirms the statements of Appendix D: Among non-uncertainty strategies, only GEEM can consistently outperform random queries. In turn, most US approaches consistently underperform random queries, a trend that may be indicative of improperly disentangled uncertainty. Only ensembles show small merit in most cases and, at least, do not fall short of random queries.

We supplement our findings by reporting the accuracy of each classifier after the acquisition budget is exhausted in Table 6. The corresponding rankings align well with the respective AUC scores Table 5. Both metrics are averaged over all 5 dataset splits and model initializations and we also report the standard deviations.

## G Computing Aleatoric and Total Confidence on CSBMs

Computing the aleatoric confidence is straightforward directly employing Definition 2.3 using Bayes rule and the generative process up to a tractable normalization constant.

$$\text{conf}^{\text{alea}}(i, c) \propto p(\mathbf{A}, \mathbf{X} \mid \mathbf{y}_{-i}^{\text{gt}}, \mathbf{y}_i = c) p(\mathbf{y}_i = c) \quad (6)$$

Computing the epistemic confidence, however, turns out to have exponential complexity in the number of unlabeled nodes  $|\mathcal{U} - i|$ :

$$\text{conf}^{\text{epi}}(i, c) = p(\mathbf{y}_i = c \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}}^{\text{gt}}) \quad (7)$$

$$= \sum_{\mathbf{y}_{\mathcal{U}-i}} p(\mathbf{y}_i = c = \mathbf{y}_{\mathcal{U}-1} \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}}^{\text{gt}}) \quad (8)$$

$$\propto \sum_{\mathbf{y}_{\mathcal{U}-i}} p(\mathbf{A}, \mathbf{X} \mid \mathbf{y}_i = c, \mathbf{y}_{\mathcal{U}-i}, \mathbf{y}_{\mathcal{O}}^{\text{gt}}) p(\mathbf{y}_i = c, \mathbf{y}_{\mathcal{U}-i}) \quad (9)$$

**Approximating Epistemic Confidence.** For larger graphs, this quickly becomes intractable and we therefore rely on a variational mean-field approximation of the joint distribution  $p(\mathbf{y}_{\mathcal{U}} \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}})$  to obtain marginals similar to [65, 66].

$$p(\mathbf{y}_{\mathcal{U}} \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}}) \approx q(\mathbf{y}_{\mathcal{U}}) := \prod_{i \in \mathcal{U}} q_i(\mathbf{y}_i) \quad (10)$$

Since the variational distributions  $q_i$  are discrete, we can fully describe them with parameters  $\gamma_{i,c} := q_i(\mathbf{y}_i = c)$ . The ELBO of this variational problem is given by:

$$J(\gamma) = \log p(\mathbf{X}, \mathbf{A}, \mathbf{y}_{\mathcal{O}}) + \mathbb{KL} \left[ q(\mathbf{y}_{\mathcal{U}}) \parallel p(\mathbf{y}_{\mathcal{U}} \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}}) \right] \quad (11)$$

$$= \mathbb{E}_{\mathbf{y}_{\mathcal{U}} \sim q} [\log p(\mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{U}} \mid \mathbf{y}_{\mathcal{O}})] + \mathbb{H} [q(\mathbf{y}_{\mathcal{U}})] + \text{const} \quad (12)$$

$$= \sum_{i,c} \gamma_{i,c} \log p(\mathbf{y}_i = c) + \sum_{i < j} \sum_{c,c'} \gamma_{i,c} \gamma_{j,c'} \log p(\mathbf{A}_{i,j} \mid \mathbf{y}_i = c, \mathbf{y}_j = c') \quad (13)$$

$$+ \sum_{i,c} \gamma_{i,c} \log p(\mathbf{X}_i \mid \mathbf{y}_i = c) + \text{const} \quad (14)$$

Here, we introduced  $\gamma_{j,c} = \mathbb{1}(c = \mathbf{y}_j^{\text{gt}})$  for all  $j \in \mathcal{O}$  for convenience. This gives rise to a constrained optimization problem where  $\sum_c \gamma_{i,c} = 1$  for all  $i \in \mathcal{U}$ . Solving this problem analytically gives rise to the equations:

$$\gamma_{i,c} \propto \exp \left( \log p(\mathbf{y}_i = c) + \sum_{j \neq i} \sum_{c'} \gamma_{j,c'} \log p(\mathbf{A}_{i,j} \mid \mathbf{y}_i = c, \mathbf{y}_j = c') + \log p(\mathbf{X}_i \mid \mathbf{y}_i = c) \right) \quad (15)$$

The marginal probabilities  $\gamma_{i,c}$  can then directly be optimized by the resulting fixed point iteration scheme, where after each iteration the probabilities are normalized:

$$\log \gamma_{i,c}^{(t+1)} = \log p(\mathbf{y}_i = c) + \sum_{j \neq i} \sum_{c'} \gamma_{j,c'}^{(t)} \log p(\mathbf{A}_{i,j} \mid \mathbf{y}_i = c, \mathbf{y}_j = c') + \log p(\mathbf{X}_i \mid \mathbf{y}_i = c) \quad (16)$$

The quality of the acquisition function given by the epistemic confidence relies on the quality of the approximated marginals  $p(\mathbf{y}_i \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}})$ , i.e. the total confidence  $\text{conf}^{\text{total}}(i, \cdot)$ . To verify

that the proposed variational scheme indeed provides reasonable approximations, we report the absolute approximation errors  $|q(\mathbf{y}_i = c) - p(\mathbf{y}_i = c \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_\mathcal{O})|$  for CSBM graphs such that exact computation of the true marginals is tractable, i.e. up to 12 nodes.

Figure 6 shows the distribution of approximation errors averaged over five different samples from a CSBM generative distribution with a structural SNR of  $\sigma_A = 2.0$  and a feature SNR  $\sigma_X = 1.0$  as well as an expected node degree  $\mathbb{E}[\deg(v)] = 4.0$ . In general, we expect the approximation to be of higher quality the more decoupled the marginals  $p(\mathbf{y}_i \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_\mathcal{O})$  are [66]. While the median error does not exceed 5%, we observe some outliers for larger graphs. In such cases, the employed approximation is inaccurate. Nonetheless, in Appendix H we observe that even in the face of sometimes poor approximations, the proposed uncertainty framework achieves strong results. We suspect a stronger approximation to perform even better.

## H Uncertainty Sampling with Ground-Truth Uncertainty

Since for real-world datasets, the data generating process is not known, we perform our analysis on CSBMs defined in Appendix A and discuss a practical approximation later in Section 3. This allows us to compute the uncertainty estimates of Definitions 2.2 to 2.4 directly by evaluating the explicit joint likelihood of the generative process  $p(\mathbf{A}, \mathbf{X}, \mathbf{y})$  (see Appendix G). While the optimality of epistemic US holds for any data-generating process, we focus on CSBMs as they have been extensively studied as proxies for real data in node classification [67]. To isolate the effect of correctly disentangling uncertainty, we also assume the parameters of the underlying CSBM to be known to the Bayesian classifier  $f_\theta^*$ . Therefore, any discrepancies in US performance are purely linked to the disentanglement into aleatoric and epistemic factors.

We compare the performance of US using the proposed uncertainty measures to contemporary uncertainty estimators over 5 graphs with 100 nodes and 7 classes sampled from a CSBM distribution  $p(\mathbf{A}, \mathbf{X}, \mathbf{y})$  in Appendix H.

**Observations:** In agreement with Theorem 2.6, epistemic uncertainty significantly outperforms random queries as well as aleatoric and total uncertainty which we explain formally in the following Propositions H.1 and H.2. We continue to analyze which aspects of uncertainty modelling are crucial for its successful in AL.

**Disentangling Uncertainty.** We first discuss why acquiring nodes with high total uncertainty  $u^{\text{total}}$  performs worse than isolating epistemic factors: Total uncertainty favors not only informative queries but also tends to acquire labels of nodes that are associated with a high aleatoric uncertainty  $u^{\text{alea}}$ .

**Proposition H.1.** *Total uncertainty  $u^{\text{total}}(i, \mathbf{y}_i^{\text{gt}})$  of a node  $i$  is proportional to the posterior over the unobserved true labels  $\mathbf{y}_{\mathcal{U}-i}^{\text{gt}}$  after acquiring its label  $\mathbf{y}_i^{\text{gt}}$ :*

$$u^{\text{total}}(i, \mathbf{y}_i^{\text{gt}}) \propto \mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_\mathcal{O}, \mathbf{y}_i = \mathbf{y}_i^{\text{gt}}]$$

We provide a proof of Proposition H.1 in Appendix B. Acquiring nodes with maximal total uncertainty maximizes the posterior of the remaining unlabeled set  $\mathbf{y}_{\mathcal{U}-i}$ . This is problematic as one way to increase this posterior probability is to remove an aleatorically uncertain node  $i$  from the unlabeled set. Such a query will not push the posterior of the remaining nodes in  $\mathbf{y}_{\mathcal{U}}$  towards their true labels and instead improve the posterior by excluding nodes that are inherently difficult to predict. In contrast, the epistemic acquisition evaluates the joint posterior in relation to the effect of removing node  $i$  from the unlabeled set (see Theorem 2.6). In fact, acquiring aleatorically uncertain nodes directly removes inherently ambiguous nodes from the unlabeled set.

**Proposition H.2.** *Aleatoric uncertainty  $u^{\text{alea}}(i, \mathbf{y}_i^{\text{gt}})$  of a node  $i$  is proportional to the posterior over the unobserved true labels  $\mathbf{y}_{\mathcal{U}-i}^{\text{gt}}$  without acquiring its label  $\mathbf{y}_i^{\text{gt}}$ :*

$$u^{\text{alea}}(i, \mathbf{y}_i^{\text{gt}}) \propto \mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_\mathcal{O}]$$

We provide a proof for Proposition H.2 in Appendix B. Proposition H.2 explains why we observe aleatoric US to be ineffective in Appendix H. It optimizes the posterior of the remaining labels  $\mathbf{y}_{\mathcal{U}-i}^{\text{gt}}$  without considering the acquisition of  $\mathbf{y}_i^{\text{gt}}$ . Such queries do not align with AL as they neglect the additional information obtained in each iteration. To optimize predictions on all remaining nodes it is crucial to properly disentangle uncertainty into aleatoric and epistemic components and acquire epistemically uncertain labels.



**Modelling the Data Generating Process.** We also investigate the importance of the uncertainty estimator to faithfully model the true data-generating process  $p(\mathbf{A}, \mathbf{X}, \mathbf{y})$ . To that end, we ablate our proposed uncertainty measures but only consider a Bayesian classifier that erroneously exclusively models present edges  $(i, j) \in \mathcal{E}$  while neglecting  $(i, j) \notin \mathcal{E}$ :  $\hat{p}(\mathbf{A}, \mathbf{X}, \mathbf{y}) := \prod_{i < j, (i, j) \in \mathcal{E}} p(\mathbf{A}_{i,j} | \mathbf{y}_i, \mathbf{y}_j) \prod_i p(\mathbf{X}_i) \prod p(\mathbf{y}_i)$ .

We specifically pick this inaccurate model to ignore non-existing edges because of its strong resemblance to contemporary GNN architectures used at the backbone of uncertainty estimators discussed in Appendix D. They rely on variations of the message-passing framework which propagates information exclusively along existing edges  $\mathcal{E}$ . In Appendix H, we observe that employing disentangled ground-truth uncertainty based on an inaccurate generative process neglecting non-existing edges harms US even with proper uncertainty disentanglement. Thus, our analysis reveals another potential shortcoming of contemporary uncertainty estimators for graphs: They may fail to accurately learn the underlying data-generating process and thus be incapable of assessing uncertainty faithfully.

**Pitfalls in US for Graphs.** Overall, we identify two aspects required for effective US on graphs that may explain the lackluster performance of existing estimators: Most importantly, uncertainty needs to be *disentangled into aleatoric and epistemic components*. Only epistemic US theoretically and empirically aligns well with AL. Furthermore, we observe an additional performance gain when the Bayesian classifier faithfully *models the data generating process*. Hence, we propose that uncertainty estimators should be designed with special care for isolating epistemic uncertainty while at the same time being expressive enough to capture the unknown underlying data-generating process.

Also looking at larger graph sizes ( $n = 1000$ ), we find epistemic US to be the strongest approach in terms of the Bayesian classifier while total uncertainty can not match its performance but also outperforms random sampling. We also observe a decrease in performance when the generative process is modelled incorrectly. Also contemporary estimators fail to outperform random queries, confirming the results reported in Section 2 for larger graph sizes.

### H.1 Ablation of the Proposed Acquisition Strategy on Different CSBM Configurations

We also ablate the proposed acquisition strategy on different configurations of the CSBM. That is, we sample graphs from a CSBM with homogeneous and symmetric affiliation matrices and vary both the structural SNR  $\frac{p}{q}$  as well as the feature SNR  $\frac{\delta_X}{\sigma_X}$ . The underlying CSBM has 100 nodes and 4 classes and labels are sampled from a uniform prior (see Appendix E.3).

We perform AL on five graphs sampled from each configuration independently and measure the absolute improvement Uncertainty Sampling achieves in terms of AUC and accuracy after a budget of  $5C = 20$  is exhausted. In Figure 8a, we show the average improvement in AUC for each SNR configuration. We also supply a similar visualization for the accuracy after the budget is exhausted in Figure 8b.


We find that our approach achieves the strongest improvement in both metrics when the structural SNR  $\sigma_A$  is neither too high nor too small: Large values make the classification problem too easy and hence AL strategies do not have to carefully query informative nodes as the overall performance is strong even in very label scarce regimes. Interestingly, our ground-truth uncertainty estimator shows strong merit when the node features are noisy, indicating that it is crucial to pick structurally informative nodes in the graph in these regimes.

At the same time, when the structural SNR is too low (in particular drops below 1.0), our method fails to outperform random acquisition: We attribute this to the mean-field approximation to the total confidence  $\text{conf}^{\text{total}}$  described in Appendix G: We observed that for graphs with low structural SNR, the fixed point iteration of Equation (16) does not converge. Hence, the approximated marginal probabilities are poor and both the prediction as well as the uncertainty estimation based on it deteriorate.


## I Visualization of Ground-Truth Uncertainty on a Toy Example

In the following, we illustrate the behavior of this acquisition function on a small CSBM graph. Figure 9a shows a sample graph with three classes and nine nodes. The greyed-out histograms represent the aleatoric confidence. The distributions correspond to the total confidence, and the size

of the nodes indicates the inverse of the ratio of both, i.e., the epistemic uncertainty, which guides our acquisition. We note that nodes 7 and 8 are the most promising candidates, as their aleatoric prediction is confident, and in the initial step since all nodes are unlabeled, the total confidence is uniform across all nodes due to the homogenous structure of the affiliation matrix.

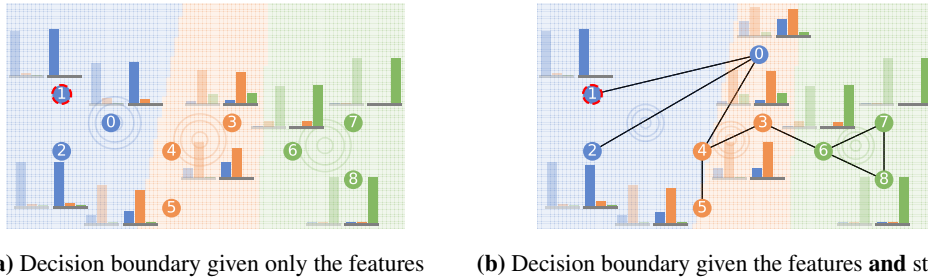


**Figure 9:** Example of CSBM with 3 classes and 9 nodes. **Left**, we can see the initial state without any labeled node. On the **right**, we can see the state after we acquired node 7.



**Figure 10:** Aleatoric uncertainty is different for symmetric nodes when the affiliation matrix is not symmetric itself.

Figure 9b depicts the subsequent iteration following the acquisition of node 7. Due to the additional information introduced, node 8 exhibits reduced epistemic uncertainty. Consequently, the most promising nodes to consider next are nodes 0 and 4, as their total confidence remains relatively low, and they display higher aleatoric confidence than their neighbors. This is because both nodes connect to two nodes from the same class, bolstering their confidence levels.



**Figure 11:** Example of CSBM with one labeled node (1). The shaded area represents the decision of the classifier. **Left**, only using the feature information. **Right** feature and structural information.

Figure 11a presents predictions made solely on feature information, with shaded regions and concentric circles demonstrating decision boundaries and class feature distributions, respectively. It is evident that incorporating feature information significantly increases confidence. Then, in Figure 11b, we adjust the position of node 0 away from the blue prediction center and introduce the structural information. The shaded regions in this figure represent the prediction for node 0 and a noticeable shift to the right

compared to Figure 11a can be observed. This shift demonstrates the impact of structural information on the classifier’s decision boundary.

Finally, Figure 10 underscores the significance of modeling non-existing edges in uncertainty estimation. Despite their symmetric neighborhoods, the aleatoric uncertainty (without feature information) differs for nodes 4 and 6 as in the underlying CSBM model the green class is less likely to not associate with the blue class. This reveals that sensible uncertainty estimators need to also utilize the *absence of edges*.

## J Approximating Disentangled Uncertainty on Real-World Data

### J.1 Multiple Pseudo-Labels (MP)

Equation (4) computes epistemic confidence as a ratio of total and aleatoric confidences. For a given classifier  $f_\theta(\mathbf{A}, \mathbf{X})$  that should be improved with AL, we interpret its predictive distribution as total confidence which is defined in Definition 2.2.

$$\text{conf}^{\text{total}}(i, c) = \mathbb{E}_{p(\theta|\mathbf{A}, \mathbf{X}, \mathbf{y}_\mathcal{O}^{\text{gt}})} [\mathbb{P}[\mathbf{y}_i = c \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_\mathcal{O} = \mathbf{y}_\mathcal{O}^{\text{gt}}, \theta]] \approx f_\theta(\mathbf{A}, \mathbf{X})_{i,c} \quad (17)$$

Note that total confidence is defined as the marginal predictive distribution associated with a node  $i$  after conditioning on the training set  $\mathbf{y}_\mathcal{O} = \mathbf{y}_\mathcal{O}^{\text{gt}}$ . This conditioning is represented by the posterior over the model parameters  $p(\theta \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_\mathcal{O}^{\text{gt}})$  that is obtained through training the classifier. In the case of deterministic GNNs, this posterior collapses to a point estimate.

Estimating the aleatoric confidence defined in Definition 2.3 is more challenging as the predictive distribution of a node  $i$  needs to be conditioned on unavailable ground-truth labels  $\mathbf{y}_{\mathcal{U}-i}^{\text{gt}}$ . We approximate those by using the predictions of  $f_\theta$  as pseudo-labels:

$$\hat{\mathbf{y}}_j = \text{argmax}_{c \in [C]} f_\theta(\mathbf{A}, \mathbf{X})_{j,c} \quad (18)$$

Similar to the approximation for total confidence, we employ the predictive distribution of a classifier to estimate aleatoric confidence. Since the set of labels that this classifier must be conditioned on the pseudo-labels  $\hat{\mathbf{y}}_{\mathcal{U}-i}$  as well, a separate classifier must be trained. For each node  $i$  that the aleatoric confidence should be predicted, we get a different set of labels the classifier needs to be trained on:  $\mathbf{y}_\mathcal{O}^{\text{gt}} \cup \hat{\mathbf{y}}_{\mathcal{U}-i}$ . This implies that  $\mathcal{O}(n)$  auxiliary classifiers need to be trained in each iteration of AL using this approximate method. For each classifier  $f_{\hat{\theta}_i}$ , we use its predictive distribution as an estimate of aleatoric confidence:

$$\text{conf}^{\text{alea}}(i, c) = \mathbb{E}_{p(\hat{\theta}_i|\mathbf{A}, \mathbf{X}, \mathbf{y}_\mathcal{O}^{\text{gt}}, \hat{\mathbf{y}}_{\mathcal{U}-i})} [\mathbb{P}[\mathbf{y}_i = c \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_\mathcal{O} = \mathbf{y}_\mathcal{O}^{\text{gt}}, \mathbf{y}_{\mathcal{U}-i} = \hat{\mathbf{y}}_{\mathcal{U}-i}, \hat{\theta}]] \approx f_{\hat{\theta}_i}(\mathbf{A}, \mathbf{X})_{i,c} \quad (19)$$

Again, the posterior  $p(\hat{\theta}_i \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_\mathcal{O}^{\text{gt}}, \hat{\mathbf{y}}_{\mathcal{U}-i})$  over the weights of the auxiliary classifier are obtained through training on both available ground-truth labels as well as pseudo labels. This implies that each auxiliary classifier  $f_{\hat{\theta}_i}$  is trained on  $n - 1$  labeled nodes. To get a scalar estimate of epistemic uncertainty that serves as an acquisition proxy, we take the ratio of both approximations according to Equation (4) and evaluate it at  $c = \hat{\mathbf{y}}_i$  as the true class label of node  $i$  is not available:

$$\text{conf}^{\text{epi}}(i, \hat{\mathbf{y}}_i) = \frac{\text{conf}^{\text{alea}}(i, \hat{\mathbf{y}}_i)}{\text{conf}^{\text{total}}(i, \hat{\mathbf{y}}_i)} \quad (20)$$

When employing the MP approximation as a strategy, in each iteration, we query the label of node  $i^* = \text{argmax}_{i \in \mathcal{U}} \text{conf}^{\text{epi}}(i, \hat{\mathbf{y}}_i)$  that maximizes (approximate) epistemic uncertainty.

### J.2 Expected Single Pseudo-Label (ESP)

An alternative approach is to not approximate epistemic uncertainty as a ratio of total and aleatoric confidences but directly estimate the right-hand side of Theorem 2.6 which we prove to be equivalent.

Both in the numerator and denominator, we need to compute joint probabilities over  $\mathbf{y}_{\mathcal{U}-i}$  which we approximate as a product of marginal probabilities predicted by the classifier. To that end, the joint probability in the denominator is conditioned on  $\mathbf{y}_{\mathcal{O}}$  only, so we again can make use of the predictive distribution of the classifier  $f_{\theta}$  to be trained. As the denominator of Theorem 2.6 requires evaluating the joint probability at the unavailable ground-truth labels  $\mathbf{y}_{\mathcal{U}-i}^{\text{gt}}$ , we again have to rely on the pseudo-labels  $\hat{\mathbf{y}}$  predicted by  $f_{\theta}$ .

$$\begin{aligned}
 \mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}}] &\approx \mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \hat{\mathbf{y}}_{\mathcal{U}-i} \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}}] \\
 &\approx \prod_{j \in \mathcal{U}-i} f_{\theta}(\mathbf{A}, \mathbf{X})_{j, \hat{\mathbf{y}}_j} \\
 &= \prod_{j \in \mathcal{U}-i} f_{\theta}(\mathbf{A}, \mathbf{X})_{j, \hat{\mathbf{y}}_j} \frac{f_{\theta}(\mathbf{A}, \mathbf{X})_{i, \hat{\mathbf{y}}_i}}{f_{\theta}(\mathbf{A}, \mathbf{X})_{i, \hat{\mathbf{y}}_i}} \\
 &= \frac{\prod_{j \in \mathcal{U}} f_{\theta}(\mathbf{A}, \mathbf{X})_{j, \hat{\mathbf{y}}_j}}{f_{\theta}(\mathbf{A}, \mathbf{X})_{i, \hat{\mathbf{y}}_i}} \\
 &\propto \frac{1}{f_{\theta}(\mathbf{A}, \mathbf{X})_{i, \hat{\mathbf{y}}_i}}
 \end{aligned} \tag{21}$$

In the last line of Equation (21), we recognized that the term  $\prod_{j \in \mathcal{U}} f_{\theta}(\mathbf{A}, \mathbf{X})_{j, \hat{\mathbf{y}}_j}$  is independent of node  $i$  which epistemic uncertainty is to be approximated for. Since for AL, we query a node that maximizes epistemic uncertainty, we can discard constant multiplicative factors that are the same for the computation of each node.

For any node  $i$ , approximating the numerator of Theorem 2.6 requires conditioning on its unavailable true label  $\mathbf{y}_i^{\text{gt}}$ . One possible remedy would be to again use the pseudo-label  $\hat{\mathbf{y}}_i$  instead. However, since we need to only condition on one additional approximate label, it is also feasible to factor in the belief of  $f_{\theta}$  about this label more accurately by taking an expectation over all possible realizations  $c \in [C]$  with respect to the predictive distribution  $f_{\theta}(\mathbf{A}, \mathbf{X})_{i, :}$ . In contrast, when computing MP approximation using Equation (19), we condition on all  $\mathbf{y}_{\mathcal{U}-i}$  simultaneously: Taking an expectation with respect to the belief of the classifier  $f_{\theta}$  would thus require  $\mathcal{O}(C^{|\mathcal{U}|})$  evaluations, each of which involves training a surrogate model from scratch, which is intractable.

Similar to computing the MP approximation, we emulate conditioning on additional observations by training a separate classifier on augmented data. That is, for conditioning on the observation  $\mathbf{y}_i = c$ , we train a model  $f_{\hat{\theta}_{i,c}}$  and interpret its predictive distribution as the marginals of the joint distribution we try to approximate. Plugging everything together, we can estimate the numerator of Theorem 2.6 for a node  $i$  as follows:

$$\begin{aligned}
 \mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}}, \mathbf{y}_i = \mathbf{y}_i^{\text{gt}}] &\approx \mathbb{E}_{c \sim f_{\theta}(\mathbf{A}, \mathbf{X})_{i, :}} [\mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}}, \mathbf{y}_i = c]] \\
 &= \sum_c \mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}}, \mathbf{y}_i = c] f_{\theta}(\mathbf{A}, \mathbf{X})_{i, c} \\
 &\approx \sum_c \left( \prod_{j \in \mathcal{U}-i} f_{\hat{\theta}_{i,c}}(\mathbf{A}, \mathbf{X})_{j, \hat{\mathbf{y}}_j^{(i,c)}} \right) f_{\theta}(\mathbf{A}, \mathbf{X})_{i, c}
 \end{aligned} \tag{22}$$

Note that here we approximate  $\mathbf{y}_{\mathcal{U}-i}^{\text{gt}}$  not with the pseudo-labels of  $f_{\theta}$ , but instead use the pseudo-labels of the surrogate classifier  $f_{\hat{\theta}_{i,c}}$  which are defined as:

$$\hat{\mathbf{y}}_j^{(i,c)} := \operatorname{argmax}_{c \in [C]} f_{\hat{\theta}_{i,c}}(\mathbf{A}, \mathbf{X})_{j, c} \tag{23}$$

We can now combine both Equations (21) and (22) and compute an estimate of the epistemic uncertainty associated with a node  $i$ :

$$\begin{aligned}
 u^{\text{epi}}(i, \mathbf{y}_i^{\text{gt}}) &= \frac{\mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}}, \mathbf{y}_i = \mathbf{y}_i^{\text{gt}}]}{\mathbb{P}[\mathbf{y}_{\mathcal{U}-i} = \mathbf{y}_{\mathcal{U}-i}^{\text{gt}} \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_{\mathcal{O}}]} \\
 &\approx \frac{\sum_c \left( \prod_{j \in \mathcal{U}-i} f_{\hat{\theta}_{i,c}}(\mathbf{A}, \mathbf{X})_{j, \tilde{\mathbf{y}}_j^{(i,c)}} \right) f_{\theta}(\mathbf{A}, \mathbf{X})_{i,c}}{\frac{\prod_{j \in \mathcal{U}} f_{\theta}(\mathbf{A}, \mathbf{X})_{j, \tilde{\mathbf{y}}_j}}{f_{\theta}(\mathbf{A}, \mathbf{X})_{i, \tilde{\mathbf{y}}_i}}} \\
 &= \frac{\sum_c \left( \prod_{j \in \mathcal{U}-i} f_{\hat{\theta}_{i,c}}(\mathbf{A}, \mathbf{X})_{j, \tilde{\mathbf{y}}_j^{(i,c)}} \right) f_{\theta}(\mathbf{A}, \mathbf{X})_{i,c}}{\prod_{j \in \mathcal{U}} f_{\theta}(\mathbf{A}, \mathbf{X})_{j, \tilde{\mathbf{y}}_j}} f_{\theta}(\mathbf{A}, \mathbf{X})_{i, \tilde{\mathbf{y}}_i} \\
 &\propto \sum_c \left( \prod_{j \in \mathcal{U}-i} f_{\hat{\theta}_{i,c}}(\mathbf{A}, \mathbf{X})_{j, \tilde{\mathbf{y}}_j^{(i,c)}} \right) f_{\theta}(\mathbf{A}, \mathbf{X})_{i,c} f_{\theta}(\mathbf{A}, \mathbf{X})_{i, \tilde{\mathbf{y}}_i}
 \end{aligned} \tag{24}$$

### J.3 Implementation

**Efficiency.** Both MP and ESP algorithms require training auxiliary classifiers on augmented datasets in order to estimate probabilities conditioned on unavailable observations. In each iteration of AL, the MP approximation requires training  $|\mathcal{U}| \in \mathcal{O}(n)$  classifiers. The ESP approximation additionally takes an expectation over the unobserved class of each node, resulting in the training of  $\mathcal{O}(nc)$  additional models per query. Notably, the auxiliary models used in the MP approximation are trained on  $n - 1$  labels while the classifiers  $f_{\hat{\theta}_{i,c}}$  the RP algorithm relies are only trained on  $|\mathcal{O}| + 1$  labels. Therefore, even though the complexity of the MP is better in theory, in practice the runtime of the ESP approximation is significantly shorter than the MP paradigm. In fact, the MP approximation did not finish an AL run within 72 hours for two datasets (see Table 1).

**Backbone Architecture and Training.** Because of the aforementioned efficiency limitations, we use an SGC model as the backbone classifier for our proposed framework. Effectively, SGC is a logistic regression model fit to diffused node features  $\mathbf{X}$ . We use the SAGA solver which is efficient for larger datasets to approximate the aleatoric uncertainty of the ESP as it uses pseudo-labels of all nodes, and we rely on liblinear in all other cases. Furthermore, we account for class imbalances and use a regularization weight of  $\lambda = 1.0$ . To mimic the GNNs used by other baselines, we diffuse the node features  $\mathbf{X}$  for 2 iterations. While acquisition requires training and evaluating auxiliary models  $f_{\hat{\theta}}$  on pseudo-labels  $\hat{\mathbf{y}}$  or  $\tilde{\mathbf{y}}$  for both approximation frameworks, we only train the underlying classifier  $f_{\theta}$  that we report numbers on using ground-truth labels iteratively revealed by the oracle.

**Discussion.** We also point out reasons for the proposed method to not be at its full efficacy yet due to the various assumptions and approximations we make. Specifically, different sources of error in estimating epistemic uncertainty can stem from (i) The classifiers  $f_{\theta}$  and  $f_{\hat{\theta}_{i,c}}$  may not faithfully model the true generative process, as described in Section 2, which results in suboptimal performance. (ii) The pseudo-labels  $\hat{\mathbf{y}}$  may not match the true labels  $\mathbf{y}$ , which in turn leads to errors in approximating aleatoric confidence and querying epistemic confidence at the correct label. (iii) The classifiers may be poorly calibrated, a tendency exhibited by some GNN architectures [15]. In fact, we observed similar experiments using a GCN instead of an SGC as a backbone to be unsuccessful. While both MP and ESP approximation paradigms aim to estimate the same quantity, they rely on different assumptions and approximations. Therefore, it is expected that they behave differently in practice. Nonetheless, they both are indicative of the potential properly disentangled uncertainty brings to AL. Indeed, as Table 1 verifies even under various approximations disentangling uncertainty greatly improves US. We suspect the ESP to perform significantly better because it relies less on pseudo-labels: It factors in the belief of the classifier  $f_{\theta}$  more accurately and only trains surrogate models on one unobserved class label.

Figures 1 and 12a to 12d showcase the practical applicability of this framework. We compare both MP and ESP approximations to the best-performing US and non-US AL strategies. MP outperforms other uncertainty estimators on Citeseer, while showing worse-than-random performance on Amazon Photos. The ESP approximation consistently yields strong results that outperform other epistemic uncertainty estimators. This underlines that disentangled epistemic uncertainty, in many instances, has the potential to be an effective guide for US. Future work can build upon the results of our analysis

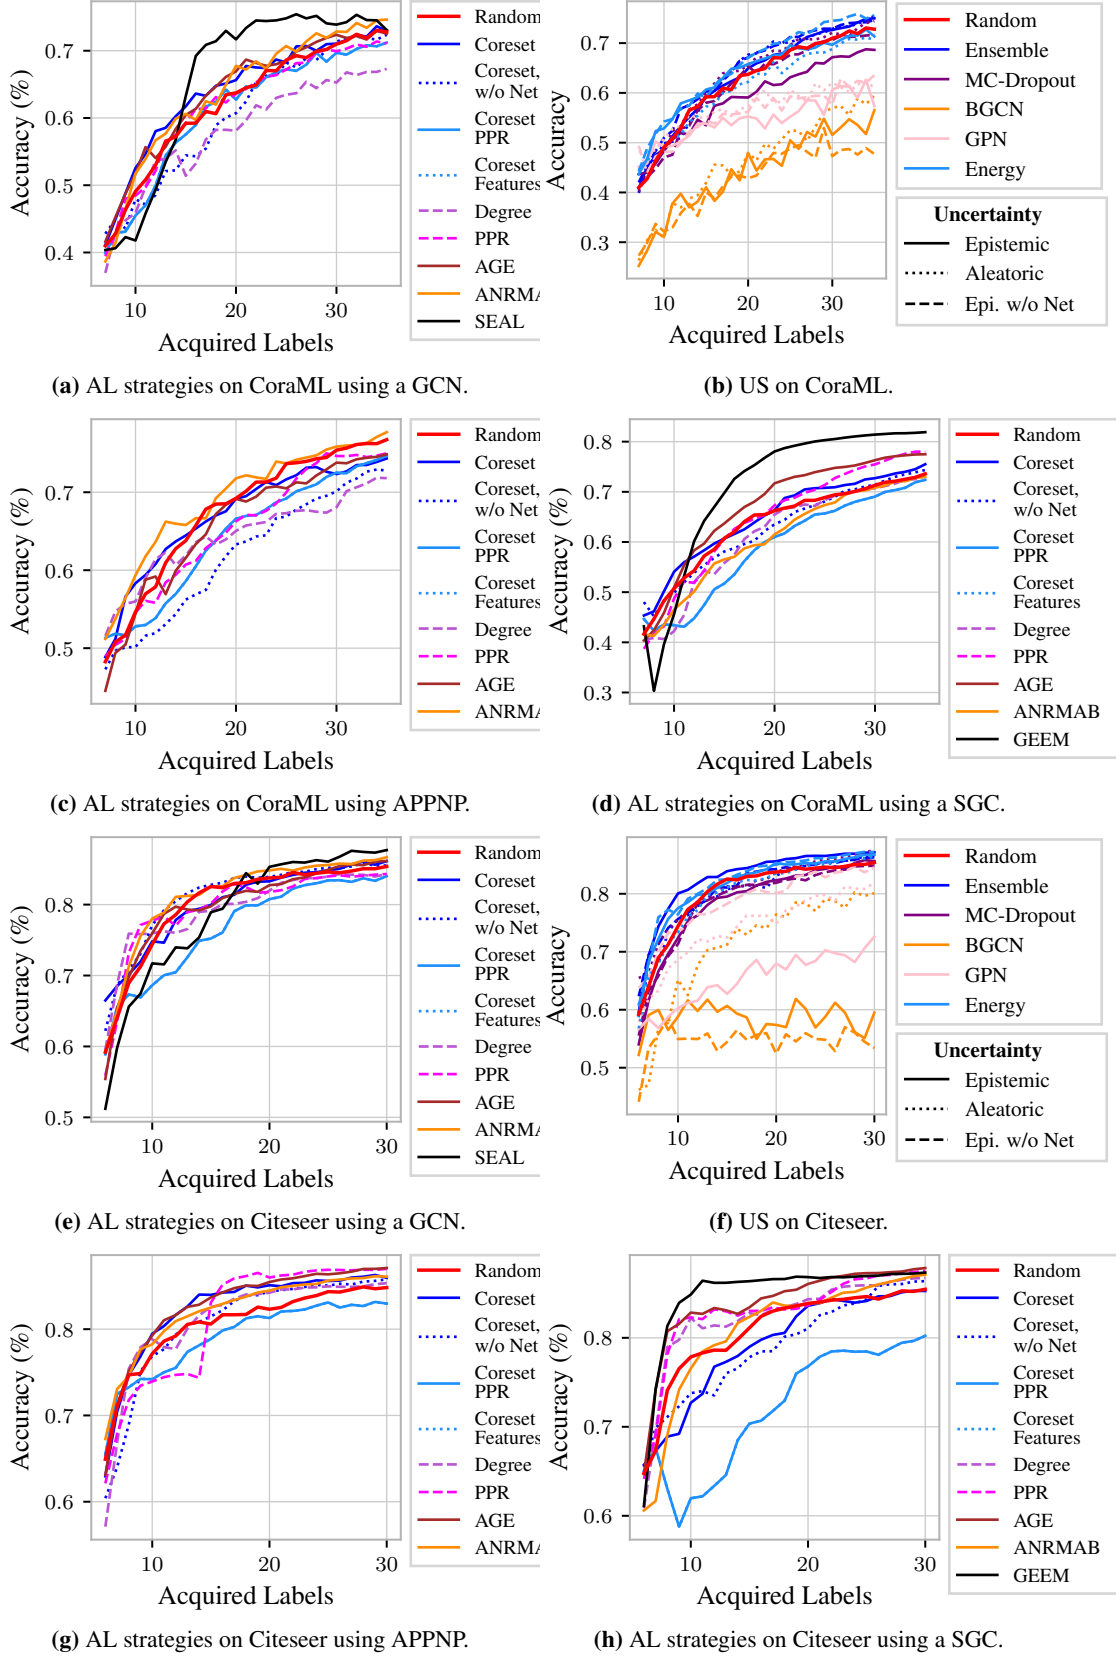


by, for example, finding more sensible and efficient approximations to uncertainty disentangling and equipping estimators with the capabilities to describe a broader family of data-generating processes.

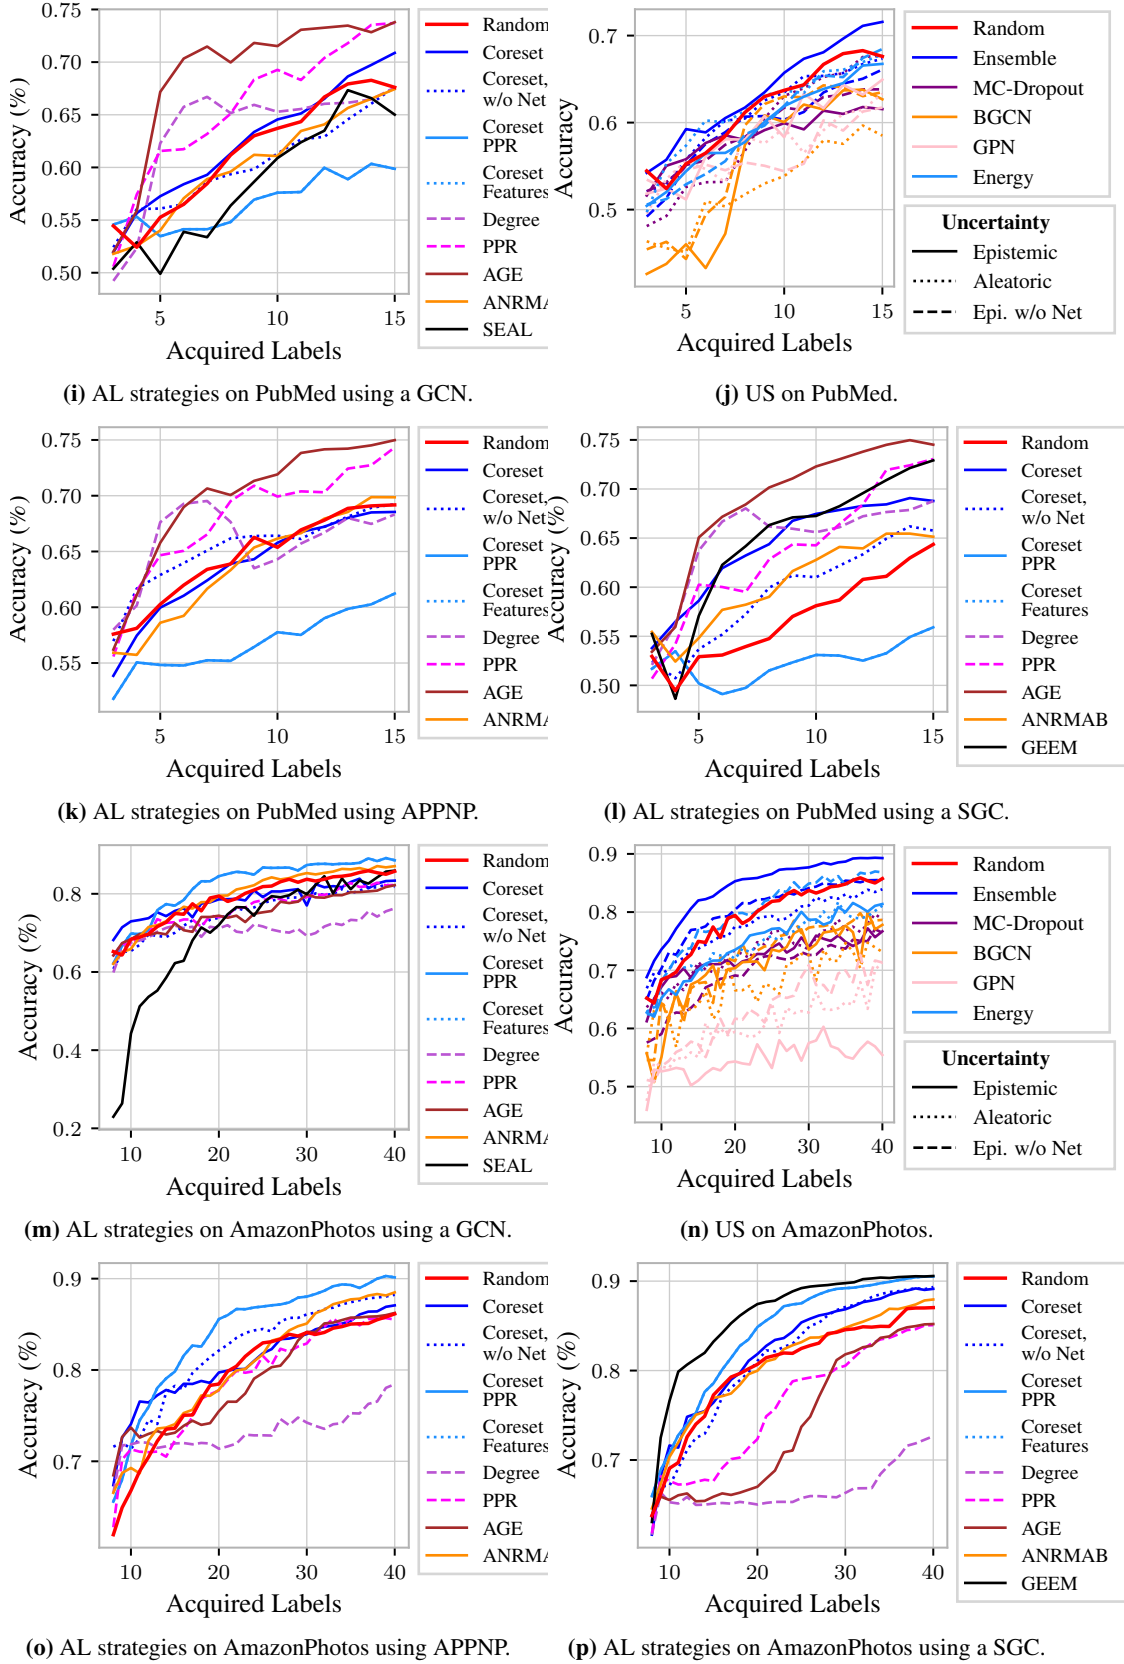
#### J.4 Approximating Ground-Truth Uncertainty without Graph Inductive Biases

We verify the impact of considering network effects when approximating ground-truth uncertainty to apply the proposed algorithm. Again, we use an SGC backbone and compute uncertainty using the aforementioned ESP approach, see Appendix J.2. We ablate estimating the marginal probabilities from the classifier that is aware of the network (by diffusing the input features) against predictions from the same model that ignores the graph. Notably, we only omit network effects when approximating uncertainty. Both training and evaluation is done with consideration for instance interdependence to enable a fair comparison between the two US strategies.

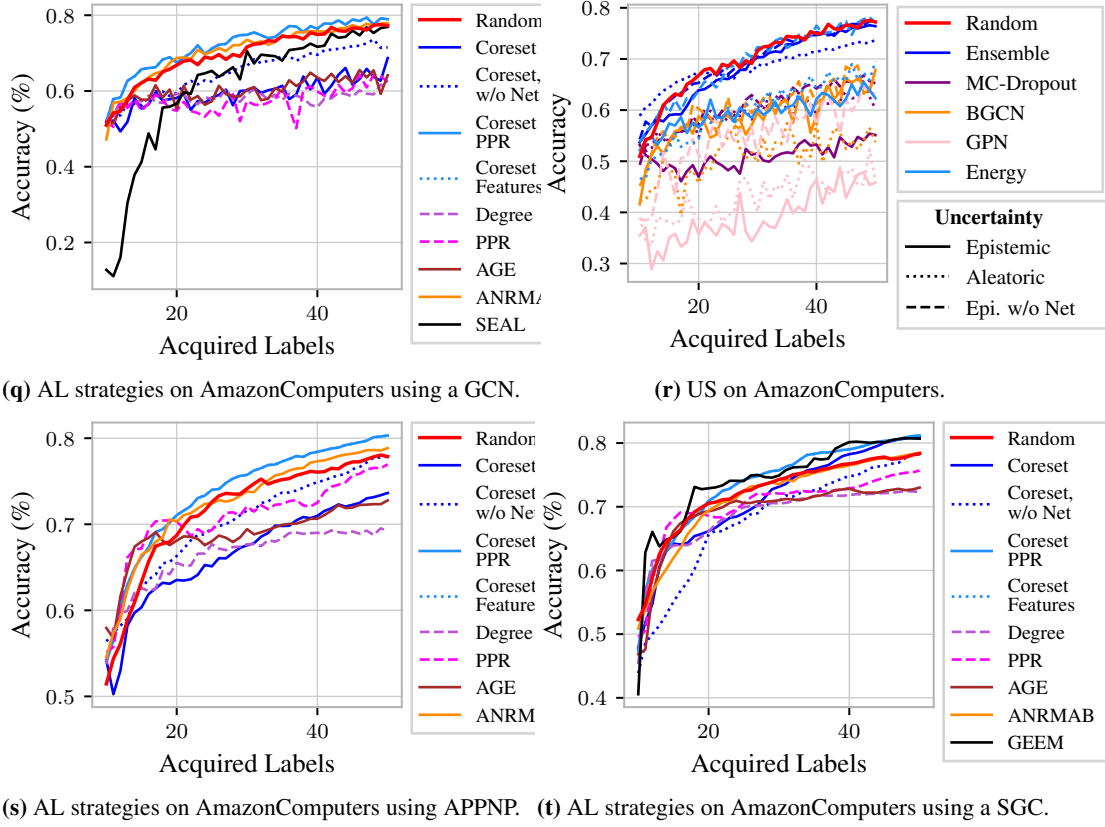
Figure 13 shows that in all settings, performance deteriorates when ignoring network effects. This is in line with the discussion in ??: The graph allows a faithful approximation of the underlying generative process and enables the ESP algorithm to identify informative queries. As has been shown by [58], modelling instance interdependence sufficiently mitigates the need for complex feature transformations. While a recent line of works regarding Energy-based models (EBMs) [68] explores how to interpret classifiers on i.i.d. data as surrogates for the data generating process, for such problems typically significantly more complex feature transformations are required. Since we can not rely on strong inductive biases implied by the graph structure, faithfully approximating the generating process and thus also ground-truth uncertainty becomes more challenging.



**Figure 4:** Accuracy curves of AL strategies, both non-uncertainty-based as well as US on different datasets for different models.



**Figure 4:** Accuracy curves of AL strategies, both non-uncertainty-based as well as US on different datasets for different models (cont.).



**Figure 4:** Accuracy curves of AL strategies, both non-uncertainty-based as well as US on different datasets for different models (cont.).

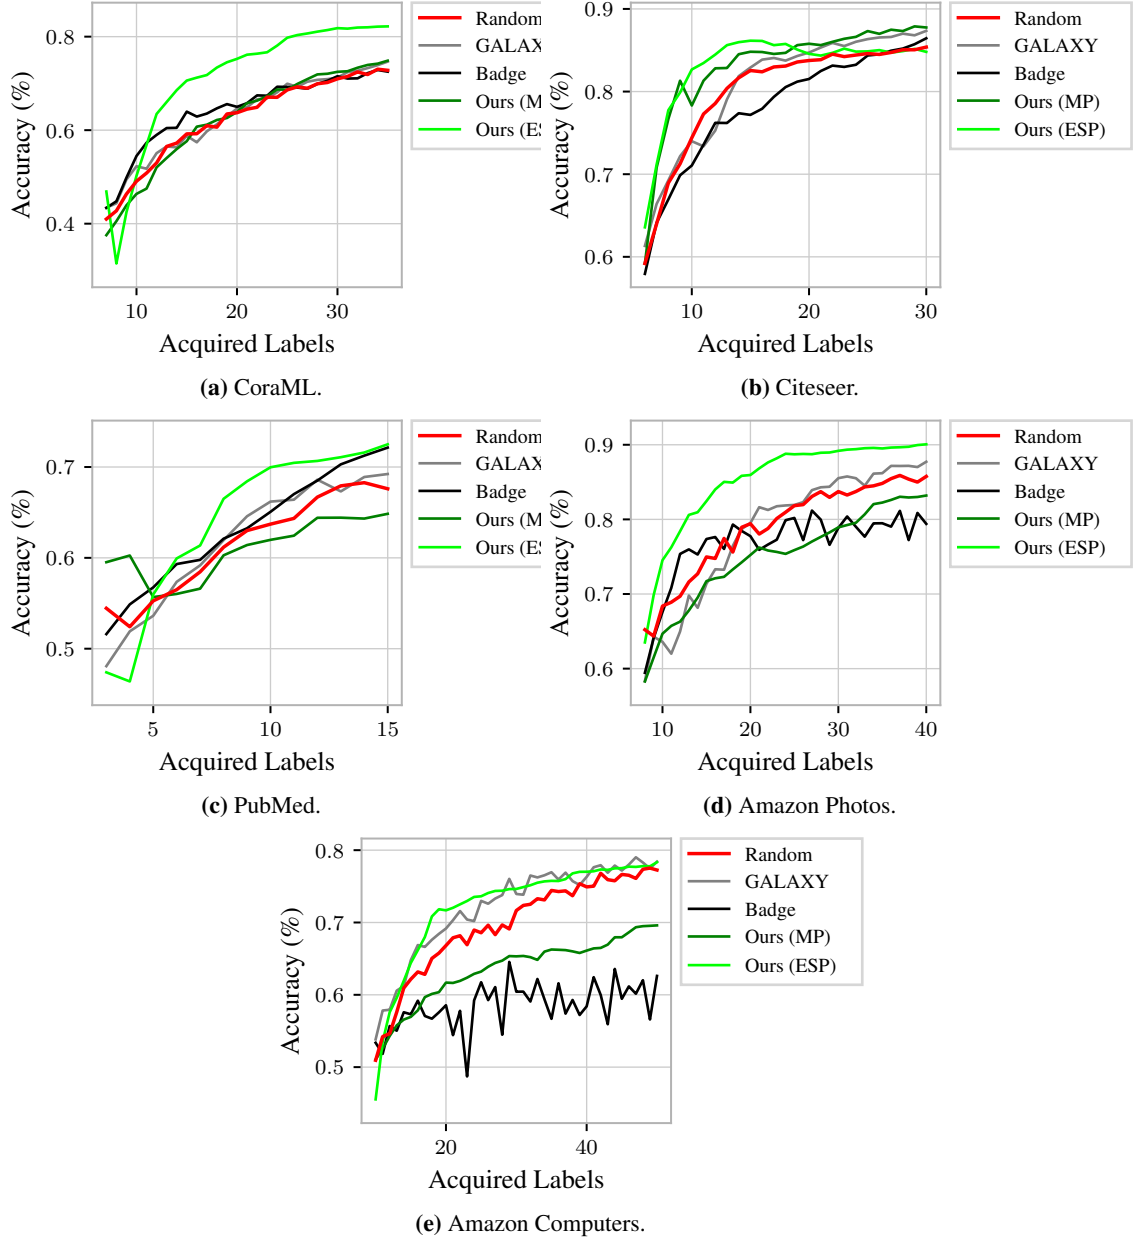
**Table 5:** Average AUC ( $\uparrow$ ) for different acquisition strategies on different models and datasets and the corresponding standard deviation over 5 different dataset splits and 5 model initializations each. We mark the best strategy per model in bold and underline the second best. For each dataset, we additionally mark the overall best model and strategy with the  $\dagger$  symbol.

	Baselines		A & X					Non-Uncertainty				X		A & X		Uncertainty		X	
	Inputs	Model	Random	CoreSet	AGE	ANRMAB	GEEM	SEAL	GALAXY	Badge	CoreSet PPR	PPR	Degree	CoreSet w/o Net	CoreSet Inputs	Epi./ (Energy)	Alca.		Epi./ (Energy)
CoramL	GCN	62.51 $\pm$ 4.31	64.35 $\pm$ 4.27	64.12 $\pm$ 3.67	64.24 $\pm$ 4.27	n/a	n/a	66.07 $\pm$ 3.94	63.24 $\pm$ 5.49	64.51 $\pm$ 4.14	59.53 $\pm$ 7.57	62.16 $\pm$ 4.53	58.39 $\pm$ 5.01	60.73 $\pm$ 4.33	61.26 $\pm$ 4.35	63.97 $\pm$ 6.29	61.30 $\pm$ 5.63	65.05 $\pm$ 4.28	64.33 $\pm$ 4.17
	APNP	67.72 $\pm$ 5.70	67.32 $\pm$ 4.20	66.12 $\pm$ 3.36	69.49 $\pm$ 5.92	n/a	n/a	n/a	n/a	n/a	71.04 $\pm$ 6.28	65.52 $\pm$ 3.48	64.57 $\pm$ 4.79	62.01 $\pm$ 5.36	64.49 $\pm$ 4.38	64.92 $\pm$ 7.02	67.68 $\pm$ 5.09	69.59 $\pm$ 4.28	66.69 $\pm$ 5.03
	Ensemble	63.89 $\pm$ 6.15	60.55 $\pm$ 5.52	64.80 $\pm$ 4.33	65.10 $\pm$ 7.29	n/a	n/a	n/a	n/a	n/a	62.65 $\pm$ 7.08	65.47 $\pm$ 4.26	62.41 $\pm$ 5.04	60.52 $\pm$ 7.37	65.07 $\pm$ 5.11	63.47 $\pm$ 6.20	64.03 $\pm$ 6.68	64.80 $\pm$ 5.14	65.82 $\pm$ 4.01
	MC-Dropout	64.94 $\pm$ 4.43	64.37 $\pm$ 5.58	64.11 $\pm$ 4.37	64.06 $\pm$ 4.94	n/a	n/a	n/a	n/a	n/a	62.92 $\pm$ 6.38	63.91 $\pm$ 5.68	61.05 $\pm$ 3.99	59.58 $\pm$ 4.39	64.35 $\pm$ 4.29	59.17 $\pm$ 5.57	63.69 $\pm$ 4.39	61.82 $\pm$ 5.19	63.87 $\pm$ 4.87
	BGCN	45.70 $\pm$ 3.26	49.37 $\pm$ 3.47	51.25 $\pm$ 4.48	47.23 $\pm$ 3.04	n/a	n/a	n/a	n/a	n/a	39.43 $\pm$ 3.87	54.64 $\pm$ 2.64	48.08 $\pm$ 3.68	43.65 $\pm$ 3.73	44.89 $\pm$ 3.15	44.42 $\pm$ 2.28	46.01 $\pm$ 3.07	42.74 $\pm$ 3.57	48.11 $\pm$ 3.16
CiteSeer	GPN	56.50 $\pm$ 4.73	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	58.04 $\pm$ 4.56	59.88 $\pm$ 3.06	58.74 $\pm$ 3.53	n/a	54.02 $\pm$ 3.83	54.79 $\pm$ 3.21	57.16 $\pm$ 3.25	55.89 $\pm$ 4.29	57.21 $\pm$ 3.12
	SGC	63.85 $\pm$ 6.05	65.23 $\pm$ 5.14	67.56 $\pm$ 3.20	61.14 $\pm$ 5.50	71.39 $\dagger$ $\pm$ 3.37	n/a	n/a	n/a	n/a	60.24 $\pm$ 7.98	65.05 $\pm$ 2.85	61.60 $\pm$ 4.34	62.94 $\pm$ 5.41	59.18 $\pm$ 5.42	67.51 $\pm$ 3.64	65.66 $\pm$ 5.70	65.05 $\pm$ 4.71	67.13 $\pm$ 4.11
	GCN	80.34 $\pm$ 3.20	80.36 $\pm$ 2.81	80.14 $\pm$ 3.00	81.59 $\pm$ 3.11	n/a	n/a	79.19 $\pm$ 4.86	81.06 $\pm$ 3.59	78.41 $\pm$ 4.78	81.51 $\pm$ 2.97	79.93 $\pm$ 3.16	79.72 $\pm$ 2.47	81.41 $\pm$ 2.28	76.87 $\pm$ 5.65	81.59 $\pm$ 3.36	80.19 $\pm$ 5.10	82.26 $\pm$ 2.42	76.79 $\pm$ 4.07
	APNP	80.77 $\pm$ 3.34	82.62 $\pm$ 2.53	83.00 $\pm$ 2.49	82.31 $\pm$ 3.10	n/a	n/a	n/a	n/a	n/a	80.70 $\pm$ 4.49	81.37 $\pm$ 4.85	81.36 $\pm$ 3.32	80.89 $\pm$ 5.26	79.20 $\pm$ 7.16	80.37 $\pm$ 6.14	77.77 $\pm$ 8.04	81.94 $\pm$ 4.95	79.15 $\pm$ 4.96
	Ensemble	81.14 $\pm$ 3.99	77.61 $\pm$ 4.28	82.37 $\pm$ 2.44	81.43 $\pm$ 3.62	n/a	n/a	n/a	n/a	n/a	81.59 $\pm$ 3.71	81.42 $\pm$ 2.61	80.31 $\pm$ 2.45	80.18 $\pm$ 4.60	79.13 $\pm$ 4.71	82.94 $\pm$ 2.72	80.26 $\pm$ 6.82	81.06 $\pm$ 4.56	77.76 $\pm$ 5.03
Pubmed	GCN	81.90 $\pm$ 3.19	78.97 $\pm$ 3.47	79.86 $\pm$ 3.31	80.46 $\pm$ 2.99	n/a	n/a	n/a	n/a	n/a	81.09 $\pm$ 3.38	79.80 $\pm$ 3.58	80.82 $\pm$ 2.85	79.40 $\pm$ 4.15	78.62 $\pm$ 4.98	80.97 $\pm$ 5.47	78.89 $\pm$ 4.73	78.89 $\pm$ 4.73	76.02 $\pm$ 4.46
	MC-Dropout	70.17 $\pm$ 3.64	73.86 $\pm$ 2.14	76.30 $\pm$ 1.42	70.70 $\pm$ 3.28	n/a	n/a	n/a	n/a	n/a	69.32 $\pm$ 3.41	73.77 $\pm$ 2.73	75.55 $\pm$ 1.47	68.59 $\pm$ 3.91	67.51 $\pm$ 3.64	58.68 $\pm$ 7.96	71.23 $\pm$ 3.80	54.85 $\pm$ 6.96	70.33 $\pm$ 4.14
	BGCN	77.07 $\pm$ 5.06	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	72.83 $\pm$ 6.88	79.52 $\pm$ 2.53	80.65 $\pm$ 1.69	n/a	71.33 $\pm$ 7.55	65.31 $\pm$ 10.21	74.28 $\pm$ 7.63	78.73 $\pm$ 6.24	77.28 $\pm$ 4.85
	GPN	81.04 $\pm$ 3.06	79.38 $\pm$ 4.44	84.21 $\pm$ 1.95	81.03 $\pm$ 3.28	85.25 $\dagger$ $\pm$ 1.80	n/a	n/a	n/a	n/a	81.59 $\pm$ 4.15	83.25 $\pm$ 2.74	82.60 $\pm$ 1.49	79.21 $\pm$ 7.73	72.04 $\pm$ 5.66	75.13 $\pm$ 9.12	78.85 $\pm$ 6.36	72.94 $\pm$ 3.35	82.90 $\pm$ 1.77
	SGC	61.56 $\pm$ 4.52	62.61 $\pm$ 5.20	69.48 $\pm$ 1.12	60.31 $\pm$ 6.68	n/a	58.62 $\pm$ 5.27	62.05 $\pm$ 5.73	63.34 $\pm$ 6.01	61.71 $\pm$ 7.24	66.07 $\pm$ 4.01	63.84 $\pm$ 5.40	60.33 $\pm$ 6.04	56.71 $\pm$ 6.44	59.64 $\pm$ 6.02	61.85 $\pm$ 5.43	59.66 $\pm$ 4.23	60.34 $\pm$ 5.98	
AmazonPhotos	APNP	64.61 $\pm$ 6.05	63.88 $\pm$ 6.47	70.18 $\dagger$ $\pm$ 4.96	63.83 $\pm$ 5.85	n/a	n/a	n/a	n/a	n/a	64.21 $\pm$ 6.62	68.24 $\pm$ 4.44	66.09 $\pm$ 4.25	65.51 $\pm$ 4.18	56.87 $\pm$ 8.10	63.09 $\pm$ 7.08	62.37 $\pm$ 5.82	62.25 $\pm$ 7.74	63.95 $\pm$ 6.34
	Ensemble	59.20 $\pm$ 6.36	64.25 $\pm$ 6.60	68.26 $\pm$ 4.28	60.40 $\pm$ 6.98	n/a	n/a	n/a	n/a	n/a	61.89 $\pm$ 5.92	63.48 $\pm$ 7.18	62.78 $\pm$ 7.21	59.19 $\pm$ 6.73	56.37 $\pm$ 5.41	63.70 $\pm$ 6.26	61.37 $\pm$ 4.97	59.71 $\pm$ 8.14	61.15 $\pm$ 4.90
	MC-Dropout	58.30 $\pm$ 6.21	62.97 $\pm$ 6.29	65.24 $\pm$ 6.30	60.50 $\pm$ 4.08	n/a	n/a	n/a	n/a	n/a	61.43 $\pm$ 5.41	65.72 $\pm$ 4.92	65.29 $\pm$ 5.23	57.71 $\pm$ 5.70	56.01 $\pm$ 6.09	58.67 $\pm$ 7.52	59.07 $\pm$ 7.03	59.23 $\pm$ 6.99	62.22 $\pm$ 4.68
	BGCN	53.59 $\pm$ 5.46	59.29 $\pm$ 4.21	56.93 $\pm$ 3.14	52.68 $\pm$ 2.92	n/a	n/a	n/a	n/a	n/a	53.40 $\pm$ 4.46	59.12 $\pm$ 2.49	55.93 $\pm$ 5.47	53.63 $\pm$ 4.50	51.40 $\pm$ 4.15	55.19 $\pm$ 3.29	52.81 $\pm$ 4.77	57.09 $\pm$ 3.93	54.62 $\pm$ 5.29
	GPN	59.70 $\pm$ 6.46	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	62.08 $\pm$ 6.79	64.64 $\pm$ 3.89	58.13 $\pm$ 4.94	n/a	54.34 $\pm$ 7.91	58.82 $\pm$ 8.90	57.24 $\pm$ 8.31	56.25 $\pm$ 4.23	59.37 $\pm$ 6.96
AmazonComputers	GCN	56.79 $\pm$ 6.28	64.48 $\pm$ 4.35	69.20 $\pm$ 4.43	60.49 $\pm$ 6.30	64.82 $\pm$ 3.68	n/a	n/a	n/a	n/a	62.15 $\pm$ 5.47	63.88 $\pm$ 4.05	55.14 $\pm$ 4.69	59.57 $\pm$ 4.45	52.25 $\pm$ 6.89	62.04 $\pm$ 5.24	61.55 $\pm$ 6.05	61.04 $\pm$ 6.84	60.74 $\pm$ 6.28
	APNP	79.08 $\pm$ 3.86	78.58 $\pm$ 2.44	75.17 $\pm$ 5.50	79.97 $\pm$ 4.20	n/a	71.16 $\pm$ 3.92	78.87 $\pm$ 4.84	76.97 $\pm$ 5.12	70.20 $\pm$ 8.55	76.30 $\pm$ 3.13	71.00 $\pm$ 5.92	75.40 $\pm$ 4.41	82.71 $\pm$ 3.65	74.68 $\pm$ 7.12	74.61 $\pm$ 5.99	79.96 $\pm$ 4.84	79.63 $\pm$ 4.59	
	Ensemble	79.29 $\pm$ 6.17	81.04 $\pm$ 2.86	79.02 $\pm$ 5.26	80.35 $\pm$ 4.96	n/a	n/a	n/a	n/a	n/a	70.37 $\pm$ 7.76	79.09 $\pm$ 3.83	73.25 $\pm$ 8.33	82.31 $\pm$ 3.29	84.24 $\pm$ 4.37	79.72 $\pm$ 6.19	77.49 $\pm$ 6.13	80.46 $\pm$ 4.97	77.69 $\pm$ 5.83
	MC-Dropout	82.23 $\pm$ 2.91	80.44 $\pm$ 4.48	77.45 $\pm$ 3.41	82.77 $\pm$ 3.62	n/a	n/a	n/a	n/a	n/a	74.93 $\pm$ 8.07	77.32 $\pm$ 4.53	75.17 $\pm$ 6.07	76.95 $\pm$ 5.38	84.01 $\pm$ 3.15	77.85 $\pm$ 6.69	80.50 $\pm$ 4.83	81.25 $\pm$ 4.27	78.80 $\pm$ 4.17
	BGCN	71.22 $\pm$ 3.86	76.63 $\pm$ 4.47	74.75 $\pm$ 3.45	80.21 $\pm$ 3.25	n/a	n/a	n/a	n/a	n/a	75.32 $\pm$ 6.70	74.33 $\pm$ 3.46	69.03 $\pm$ 6.60	75.71 $\pm$ 4.62	82.45 $\pm$ 2.84	72.42 $\pm$ 6.79	73.16 $\pm$ 6.19	69.68 $\pm$ 4.19	78.80 $\pm$ 4.17
AmazonReviews	GCN	62.80 $\pm$ 4.22	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	59.34 $\pm$ 7.49	64.51 $\pm$ 4.28	61.23 $\pm$ 5.84	69.78 $\pm$ 4.24	73.39 $\pm$ 4.16	70.83 $\pm$ 3.36	67.83 $\pm$ 5.83	72.21 $\pm$ 5.84	69.19 $\pm$ 5.80
	GPN	62.80 $\pm$ 4.22	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	55.59 $\pm$ 4.09	62.17 $\pm$ 4.08	56.77 $\pm$ 4.94	n/a	65.07 $\pm$ 2.74	54.78 $\pm$ 3.57	60.53 $\pm$ 4.17	62.90 $\pm$ 2.61	62.41 $\pm$ 2.75
	SGC	80.52 $\pm$ 4.89	82.32 $\pm$ 2.58	74.01 $\pm$ 5.96	80.92 $\pm$ 4.04	86.43 $\dagger$ $\pm$ 4.13	n/a	n/a	n/a	n/a	66.94 $\pm$ 7.86	76.16 $\pm$ 4.06	66.63 $\pm$ 8.12	81.47 $\pm$ 5.08	84.24 $\pm$ 3.59	84.01 $\pm$ 4.72	71.43 $\pm$ 3.37	80.75 $\pm$ 4.23	76.38 $\pm$ 7.40
	GCN	69.80 $\pm$ 4.83	59.36 $\pm$ 4.48	69.07 $\pm$ 6.85	70.70 $\pm$ 2.66	n/a	61.51 $\pm$ 4.11	72.32 $\pm$ 2.47	58.48 $\pm$ 4.40	61.22 $\pm$ 5.64	57.40 $\pm$ 4.48	58.14 $\pm$ 5.29	65.13 $\pm$ 4.31	72.34 $\pm$ 2.75	59.62 $\pm$ 7.36	60.17 $\pm$ 6.98	69.34 $\pm$ 4.75	69.87 $\pm$ 3.77	
	APNP	71.69 $\pm$ 3.34	66.62 $\pm$ 3.55	68.89 $\pm$ 3.32	72.69 $\pm$ 3.27	n/a	n/a	n/a	n/a	n/a	65.75 $\pm$ 6.76	70.91 $\pm$ 5.92	66.42 $\pm$ 4.73	70.22 $\pm$ 4.44	73.83 $\pm$ 3.66	62.03 $\pm$ 8.44	62.26 $\pm$ 10.26	68.79 $\pm$ 5.46	71.72 $\pm$ 3.79
AmazonMusic	Ensemble	72.50 $\pm$ 3.96	64.19 $\pm$ 4.49	60.69 $\pm$ 5.43	72.73 $\pm$ 4.16	n/a	n/a	n/a	n/a	n/a	64.13 $\pm$ 6.76	66.50 $\pm$ 7.02	63.16 $\pm$ 5.33	67.20 $\pm$ 6.40	75.39 $\dagger$ $\pm$ 2.88	68.38 $\pm$ 5.04	68.47 $\pm$ 7.74	69.40 $\pm$ 4.35	73.67 $\pm$ 3.30
	MC-Dropout	68.06 $\pm$ 4.05	58.58 $\pm$ 4.30	57.04 $\pm$ 4.61	69.01 $\pm$ 2.92	n/a	n/a	n/a	n/a	n/a	64.43 $\pm$ 6.67	55.88 $\pm$ 4.81	56.72 $\pm$ 6.12	58.65 $\pm$ 3.16	71.05 $\pm$ 2.88	51.02 $\pm$ 9.05	59.31 $\pm$ 7.96	60.71 $\pm$ 7.40	70.66 $\pm$ 3.97
	BGCN	60.52 $\pm$ 3.72	45.65 $\pm$ 3.02	46.72 $\pm$ 3.74	60.32 $\pm$ 4.31	n/a	n/a	n/a	n/a	n/a	39.86 $\pm$ 8.46	43.61 $\pm$ 4.17	45.52 $\pm$ 2.93	58.85 $\pm$ 3.41	60.79 $\pm$ 3.89	58.64 $\pm$ 5.96	51.11 $\pm$ 6.23	60.19 $\pm$ 3.82	59.20 $\pm$ 3.67
	GPN	52.26 $\pm$ 5.90	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	34.71 $\pm$ 4.90	49.30 $\pm$ 5.60	45.21 $\pm$ 5.32	n/a	56.32 $\pm$ 3.79	39.21 $\pm$ 3.92	42.95 $\pm$ 4.53	54.83 $\pm$ 4.84	53.71 $\pm$ 3.91
	SGC	72.30 $\pm$ 2.74	71.53 $\pm$ 3.86	69.31 $\pm$ 4.37	71.62 $\pm$ 3.79	74.49 $\dagger$ $\pm$ 4.43	n/a	n/a	n/a	n/a	58.52 $\pm$ 6.44	70.35 $\pm$ 3.84	68.70 $\pm$ 3.24	68.14 $\pm$ 5.46	73.91 $\pm$ 3.30	61.02 $\pm$ 8.74	59.34 $\pm$ 10.50	66.53 $\pm$ 5.88	69.24 $\pm$ 5.72




**Table 6:** Average final classification accuracy ( $\uparrow$ ) for different acquisition strategies on different models and datasets and the corresponding standard deviation over 5 different dataset splits and 5 model initializations each. We mark the best strategy per model in bold and underline the second best. For each dataset, we additionally mark the overall best model and strategy with the  $\dagger$  symbol.

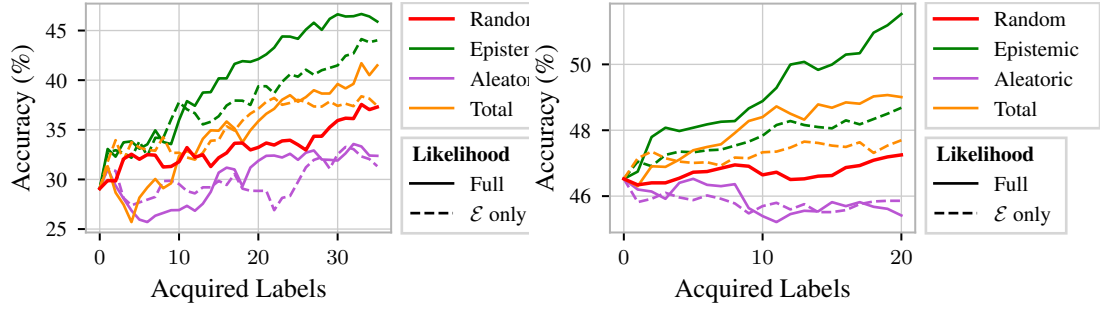
Inputs	Model	Baselines		Non-Uncertainty										Uncertainty			
		Random	Coreset	AGE	ANRMAB	GEEM	SEAL	GALAXY	BADGE	Coreset PPR	Degree	Coreset w/o Net	Coreset Inputs	Epi/(Energy)	A & X	Epi/(Energy)	X
CoraML	GCN	72.80 $\pm$ 3.74	72.95 $\pm$ 4.17	72.54 $\pm$ 3.61	74.64 $\pm$ 3.46	n/a	73.08 $\pm$ 3.13	<u>74.07</u> $\pm$ 4.82	72.51 $\pm$ 4.05	69.33 $\pm$ 6.04	71.10 $\pm$ 4.26	72.36 $\pm$ 4.27	71.21 $\pm$ 4.18	71.33 $\pm$ 5.39	70.69 $\pm$ 7.56	<b>75.69</b> $\pm$ 3.82	74.08 $\pm$ 3.79
	APPNP	76.74 $\pm$ 4.45	74.33 $\pm$ 4.02	74.86 $\pm$ 2.76	77.71 $\pm$ 4.59	n/a	n/a	n/a	n/a	76.11 $\pm$ 4.08	74.99 $\pm$ 2.24	71.81 $\pm$ 3.78	72.78 $\pm$ 3.84	72.29 $\pm$ 3.83	75.61 $\pm$ 3.88	<b>77.91</b> $\pm$ 3.03	75.29 $\pm$ 3.71
	Ensemble	74.25 $\pm$ 4.81	73.82 $\pm$ 4.05	73.82 $\pm$ 4.05	75.40 $\pm$ 4.56	n/a	n/a	n/a	n/a	72.70 $\pm$ 4.57	72.96 $\pm$ 2.64	69.77 $\pm$ 5.69	74.12 $\pm$ 3.77	74.98 $\pm$ 5.58	75.51 $\pm$ 3.88	74.63 $\pm$ 3.64	<b>75.62</b> $\pm$ 2.29
	MC-Dropout	73.80 $\pm$ 4.46	72.44 $\pm$ 5.59	72.40 $\pm$ 4.09	73.44 $\pm$ 5.36	n/a	n/a	n/a	n/a	70.91 $\pm$ 5.10	72.06 $\pm$ 5.10	71.42 $\pm$ 5.19	74.06 $\pm$ 5.19	68.64 $\pm$ 5.57	74.42 $\pm$ 4.15	71.67 $\pm$ 4.63	<b>75.86</b> $\pm$ 4.55
	BGCN	59.09 $\pm$ 8.86	61.08 $\pm$ 5.05	61.20 $\pm$ 8.35	57.80 $\pm$ 6.47	n/a	n/a	n/a	n/a	46.40 $\pm$ 5.81	<b>64.79</b> $\pm$ 4.96	56.22 $\pm$ 4.89	55.73 $\pm$ 7.12	56.53 $\pm$ 5.74	59.48 $\pm$ 7.56	47.72 $\pm$ 4.13	57.87 $\pm$ 7.44
Citeseer	GCN	64.24 $\pm$ 5.57	n/a	n/a	n/a	n/a	n/a	n/a	n/a	59.66 $\pm$ 7.23	<b>66.85</b> $\pm$ 4.48	n/a	n/a	60.52 $\pm$ 6.49	61.08 $\pm$ 7.07	65.97 $\pm$ 7.69	65.97 $\pm$ 7.69
	APPNP	73.59 $\pm$ 9.90	75.50 $\pm$ 7.70	77.49 $\pm$ 1.44	72.99 $\pm$ 4.52	<b>81.89</b> $\pm$ 1.58	n/a	n/a	n/a	66.40 $\pm$ 4.16	73.02 $\pm$ 2.26	74.35 $\pm$ 4.01	72.39 $\pm$ 4.66	77.90 $\pm$ 2.40	75.23 $\pm$ 3.96	76.59 $\pm$ 2.79	77.03 $\pm$ 3.50
	Ensemble	85.38 $\pm$ 2.24	86.13 $\pm$ 1.80	86.15 $\pm$ 1.99	86.08 $\pm$ 1.89	n/a	n/a	87.35 $\pm$ 1.64	86.45 $\pm$ 1.39	86.03 $\pm$ 1.72	84.31 $\pm$ 1.58	85.51 $\pm$ 1.90	84.01 $\pm$ 1.66	86.71 $\pm$ 2.45	86.39 $\pm$ 2.89	87.21 $\pm$ 1.98	87.39 $\pm$ 1.95
	MC-Dropout	84.82 $\pm$ 3.78	85.90 $\pm$ 2.64	87.09 $\pm$ 1.55	86.09 $\pm$ 2.08	n/a	n/a	n/a	n/a	85.31 $\pm$ 3.32	86.39 $\pm$ 1.48	85.33 $\pm$ 1.68	82.97 $\pm$ 5.03	86.59 $\pm$ 3.81	84.49 $\pm$ 5.34	85.83 $\pm$ 3.38	<b>87.15</b> $\pm$ 2.90
	BGCN	85.63 $\pm$ 3.26	85.00 $\pm$ 2.65	86.07 $\pm$ 2.35	86.77 $\pm$ 2.16	n/a	n/a	n/a	n/a	86.19 $\pm$ 3.25	85.07 $\pm$ 1.50	84.90 $\pm$ 1.59	83.98 $\pm$ 3.30	<b>87.19</b> $\pm$ 1.54	86.41 $\pm$ 1.64	86.27 $\pm$ 2.37	<b>88.06</b> $\pm$ 2.00
Pubmed	GCN	86.64 $\pm$ 1.90	86.52 $\pm$ 1.88	85.50 $\pm$ 1.83	86.07 $\pm$ 2.77	n/a	n/a	n/a	n/a	86.01 $\pm$ 2.24	85.23 $\pm$ 1.13	85.12 $\pm$ 2.47	84.74 $\pm$ 3.07	85.63 $\pm$ 2.50	87.59 $\pm$ 2.09	84.84 $\pm$ 3.57	<b>88.06</b> $\pm$ 2.00
	APPNP	78.89 $\pm$ 7.20	80.45 $\pm$ 3.60	79.59 $\pm$ 4.58	79.37 $\pm$ 4.66	n/a	n/a	n/a	n/a	75.14 $\pm$ 6.50	80.44 $\pm$ 4.52	76.46 $\pm$ 5.00	76.41 $\pm$ 3.51	59.50 $\pm$ 13.96	80.08 $\pm$ 5.25	53.42 $\pm$ 12.61	78.44 $\pm$ 5.55
	Ensemble	80.96 $\pm$ 4.76	n/a	n/a	n/a	n/a	n/a	n/a	n/a	77.06 $\pm$ 7.62	83.35 $\pm$ 2.24	n/a	n/a	72.61 $\pm$ 7.77	81.87 $\pm$ 5.27	<b>84.58</b> $\pm$ 3.86	81.60 $\pm$ 3.65
	MC-Dropout	85.45 $\pm$ 3.72	85.20 $\pm$ 3.90	<b>87.84</b> $\pm$ 1.10	87.09 $\pm$ 1.86	87.33 $\pm$ 1.95	n/a	n/a	n/a	87.17 $\pm$ 2.14	86.73 $\pm$ 1.73	86.37 $\pm$ 2.99	80.23 $\pm$ 4.86	75.63 $\pm$ 10.15	86.19 $\pm$ 3.54	80.16 $\pm$ 10.35	87.78 $\pm$ 1.98
	BGCN	67.60 $\pm$ 3.10	70.87 $\pm$ 5.94	<b>73.79</b> $\pm$ 4.08	67.44 $\pm$ 8.55	n/a	n/a	69.23 $\pm$ 5.44	72.14 $\pm$ 5.16	66.73 $\pm$ 3.80	73.72 $\pm$ 3.72	67.66 $\pm$ 5.04	59.87 $\pm$ 7.29	66.75 $\pm$ 8.22	67.84 $\pm$ 6.37	68.46 $\pm$ 7.13	65.62 $\pm$ 7.53
AmazonPhotos	GCN	69.18 $\pm$ 5.57	68.54 $\pm$ 7.37	<b>74.98</b> $\pm$ 4.48	69.86 $\pm$ 5.37	n/a	n/a	n/a	n/a	69.19 $\pm$ 5.81	74.33 $\pm$ 3.52	69.24 $\pm$ 5.44	61.23 $\pm$ 3.30	70.80 $\pm$ 5.61	68.16 $\pm$ 6.96	68.09 $\pm$ 8.57	68.93 $\pm$ 6.57
	APPNP	66.10 $\pm$ 5.96	70.89 $\pm$ 4.16	<b>72.65</b> $\pm$ 3.86	68.74 $\pm$ 6.24	n/a	n/a	n/a	n/a	68.24 $\pm$ 5.75	66.11 $\pm$ 6.92	65.03 $\pm$ 5.24	60.29 $\pm$ 6.79	71.57 $\pm$ 6.75	67.35 $\pm$ 5.07	66.10 $\pm$ 6.61	68.07 $\pm$ 5.96
	Ensemble	66.45 $\pm$ 4.99	71.19 $\pm$ 7.79	71.15 $\pm$ 6.03	65.50 $\pm$ 4.56	n/a	n/a	n/a	n/a	67.88 $\pm$ 6.79	70.11 $\pm$ 7.06	62.32 $\pm$ 7.93	60.82 $\pm$ 6.69	61.55 $\pm$ 8.42	67.70 $\pm$ 7.60	63.82 $\pm$ 7.51	69.75 $\pm$ 4.93
	MC-Dropout	62.57 $\pm$ 8.73	66.70 $\pm$ 9.98	63.77 $\pm$ 5.06	60.72 $\pm$ 6.22	n/a	n/a	n/a	n/a	59.02 $\pm$ 9.08	<b>68.97</b> $\pm$ 3.46	60.36 $\pm$ 6.70	54.49 $\pm$ 3.35	62.67 $\pm$ 6.46	58.51 $\pm$ 7.89	63.49 $\pm$ 6.15	61.25 $\pm$ 7.83
	BGCN	63.75 $\pm$ 8.43	n/a	n/a	n/a	n/a	n/a	n/a	n/a	66.96 $\pm$ 9.94	<b>67.79</b> $\pm$ 4.98	n/a	54.32 $\pm$ 9.25	64.94 $\pm$ 7.58	63.31 $\pm$ 8.19	61.74 $\pm$ 6.70	65.38 $\pm$ 9.24
AmazonComputers	GCN	64.36 $\pm$ 8.31	68.78 $\pm$ 3.81	<b>74.51</b> $\pm$ 4.79	65.14 $\pm$ 7.04	72.91 $\pm$ 5.03	n/a	n/a	n/a	67.71 $\pm$ 3.66	68.76 $\pm$ 4.42	65.76 $\pm$ 4.13	55.91 $\pm$ 8.35	70.29 $\pm$ 6.11	67.61 $\pm$ 5.85	67.19 $\pm$ 7.03	68.18 $\pm$ 6.85
	APPNP	85.76 $\pm$ 3.46	83.37 $\pm$ 3.57	82.20 $\pm$ 2.98	87.07 $\pm$ 3.40	n/a	85.92 $\pm$ 3.50	<u>87.72</u> $\pm$ 4.07	79.40 $\pm$ 7.22	75.73 $\pm$ 3.56	82.28 $\pm$ 3.15	81.76 $\pm$ 3.01	<b>88.58</b> $\pm$ 2.45	81.39 $\pm$ 3.57	81.12 $\pm$ 5.41	86.82 $\pm$ 4.51	84.73 $\pm$ 4.33
	Ensemble	86.14 $\pm$ 3.30	87.07 $\pm$ 1.60	86.24 $\pm$ 2.13	88.49 $\pm$ 2.49	n/a	n/a	n/a	n/a	81.93 $\pm$ 5.54	85.52 $\pm$ 3.33	88.24 $\pm$ 1.60	<b>90.14</b> $\pm$ 1.75	86.07 $\pm$ 4.41	84.31 $\pm$ 4.21	86.65 $\pm$ 3.87	85.28 $\pm$ 4.15
	MC-Dropout	87.97 $\pm$ 2.14	84.82 $\pm$ 3.48	83.92 $\pm$ 4.04	88.10 $\pm$ 3.30	n/a	n/a	n/a	n/a	79.82 $\pm$ 7.86	83.36 $\pm$ 4.01	83.51 $\pm$ 3.15	<b>89.96</b> $\pm$ 1.61	89.25 $\pm$ 2.36	83.92 $\pm$ 6.09	85.53 $\pm$ 4.31	87.17 $\pm$ 3.75
	BGCN	79.22 $\pm$ 7.13	69.90 $\pm$ 9.01	73.08 $\pm$ 9.53	<b>82.32</b> $\pm$ 4.67	n/a	n/a	n/a	n/a	79.07 $\pm$ 6.75	80.45 $\pm$ 9.57	82.11 $\pm$ 5.62	<b>88.89</b> $\pm$ 2.26	70.69 $\pm$ 5.01	78.69 $\pm$ 5.88	76.44 $\pm$ 8.43	87.09 $\pm$ 2.69
AmazonReviews	GCN	64.59 $\pm$ 8.87	n/a	n/a	n/a	n/a	n/a	n/a	n/a	56.95 $\pm$ 9.79	71.68 $\pm$ 3.41	59.87 $\pm$ 9.57	77.46 $\pm$ 10.16	55.44 $\pm$ 7.09	71.29 $\pm$ 9.69	78.66 $\pm$ 8.01	77.17 $\pm$ 7.59
	APPNP	80.13 $\pm$ 1.79	85.20 $\pm$ 2.79	87.95 $\pm$ 2.43	<b>90.57</b> $\pm$ 2.83	n/a	n/a	n/a	n/a	72.18 $\pm$ 8.94	85.08 $\pm$ 3.70	89.34 $\pm$ 2.13	<b>90.52</b> $\pm$ 1.51	90.22 $\pm$ 2.48	78.75 $\pm$ 7.06	85.71 $\pm$ 3.10	84.25 $\pm$ 5.51
	Ensemble	77.25 $\pm$ 3.74	68.69 $\pm$ 7.26	64.17 $\pm$ 49.53	77.93 $\pm$ 2.95	n/a	76.92 $\pm$ 3.00	<u>78.12</u> $\pm$ 2.18	62.61 $\pm$ 11.92	67.04 $\pm$ 6.67	62.32 $\pm$ 9.02	71.51 $\pm$ 4.11	<b>78.94</b> $\pm$ 2.69	62.35 $\pm$ 11.46	69.54 $\pm$ 8.69	77.08 $\pm$ 4.57	77.84 $\pm$ 3.28
	MC-Dropout	77.87 $\pm$ 3.37	73.64 $\pm$ 1.12	72.78 $\pm$ 4.66	78.86 $\pm$ 3.23	n/a	n/a	n/a	n/a	76.96 $\pm$ 3.02	69.24 $\pm$ 5.44	77.85 $\pm$ 2.75	<b>80.31</b> $\pm$ 2.32	67.25 $\pm$ 9.27	71.19 $\pm$ 10.07	78.44 $\pm$ 4.37	79.87 $\pm$ 4.42
	BGCN	77.88 $\pm$ 3.90	69.20 $\pm$ 5.50	66.15 $\pm$ 4.00	78.70 $\pm$ 4.04	n/a	n/a	n/a	n/a	70.53 $\pm$ 3.58	66.92 $\pm$ 5.77	71.42 $\pm$ 5.38	<b>80.50</b> $\pm$ 2.38	76.41 $\pm$ 3.94	73.71 $\pm$ 7.31	77.38 $\pm$ 5.33	80.38 $\pm$ 4.37
AmazonMusic	GCN	63.65 $\pm$ 3.93	63.65 $\pm$ 3.93	62.67 $\pm$ 8.51	75.07 $\pm$ 3.73	n/a	n/a	n/a	n/a	59.78 $\pm$ 11.42	58.54 $\pm$ 12.72	72.24 $\pm$ 4.90	<b>78.17</b> $\pm$ 3.71	55.12 $\pm$ 11.72	60.21 $\pm$ 11.38	65.25 $\pm$ 10.33	77.00 $\pm$ 3.59
	APPNP	63.23 $\pm$ 9.50	39.84 $\pm$ 15.30	45.25 $\pm$ 13.93	65.87 $\pm$ 9.88	n/a	n/a	n/a	n/a	35.80 $\pm$ 15.73	45.34 $\pm$ 12.66	63.64 $\pm$ 7.52	62.32 $\pm$ 13.64	<b>67.87</b> $\pm$ 6.76	53.71 $\pm$ 13.66	67.06 $\pm$ 8.51	63.21 $\pm$ 11.97
	Ensemble	59.89 $\pm$ 13.72	n/a	n/a	n/a	n/a	n/a	n/a	n/a	32.00 $\pm$ 8.75	60.63 $\pm$ 10.59	<b>64.79</b> $\pm$ 10.40	64.79 $\pm$ 10.40	45.85 $\pm$ 13.64	45.95 $\pm$ 14.32	63.25 $\pm$ 3.79	64.24 $\pm$ 6.71
	MC-Dropout	73.01 $\pm$ 4.68	73.01 $\pm$ 4.68	73.01 $\pm$ 4.68	73.01 $\pm$ 4.68	n/a	n/a	n/a	n/a	64.92 $\pm$ 0.75	75.69 $\pm$ 3.56	72.31 $\pm$ 3.54	<b>81.19</b> $\pm$ 2.25	68.70 $\pm$ 8.90	67.44 $\pm$ 10.26	73.46 $\pm$ 6.06	78.38 $\pm$ 4.05
	BGCN	78.35 $\pm$ 2.44	78.35 $\pm$ 2.44	78.35 $\pm$ 2.44	78.35 $\pm$ 2.44	n/a	n/a	n/a	n/a	78.35 $\pm$ 2.44	78.35 $\pm$ 2.44	78.35 $\pm$ 2.44	<b>81.19</b> $\pm$ 2.25	78.35 $\pm$ 2.44	78.35 $\pm$ 2.44	78.35 $\pm$ 2.44	78.35 $\pm$ 2.44



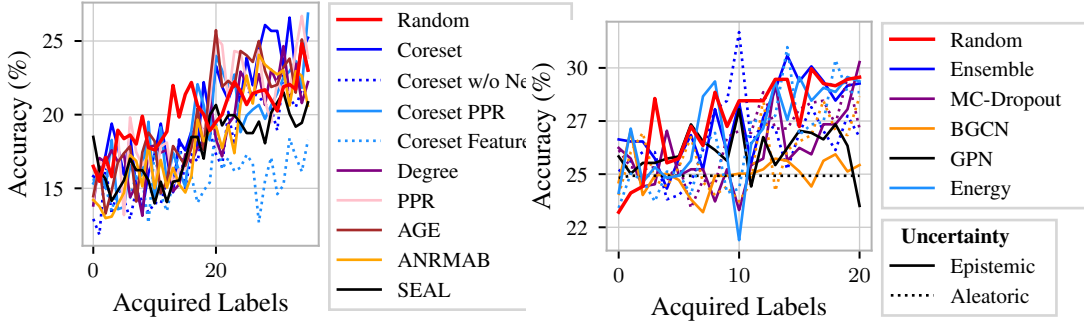
**Figure 5:** US using the i.i.d. GALAXY and BADGE baselines versus our approach (ESP) and random sampling.



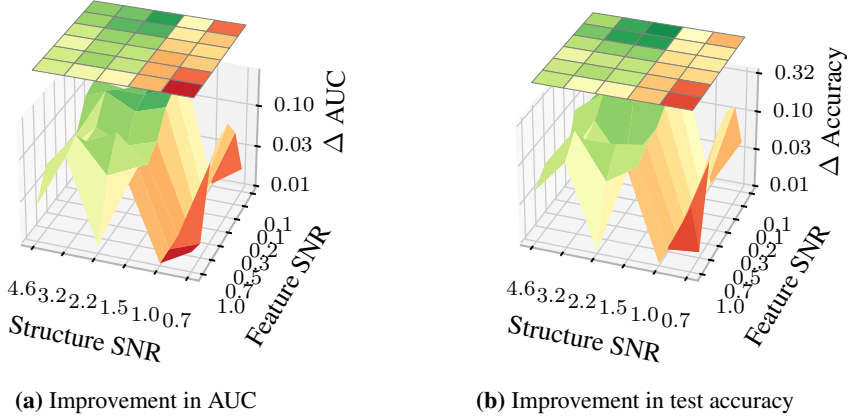
**Figure 6:** Absolute error distribution between approximate total confidence  $q(\mathbf{y}_i)$  and true total confidence  $p(\mathbf{y}_i \mid \mathbf{A}, \mathbf{X}, \mathbf{y}_O)$  for graphs of different sizes.



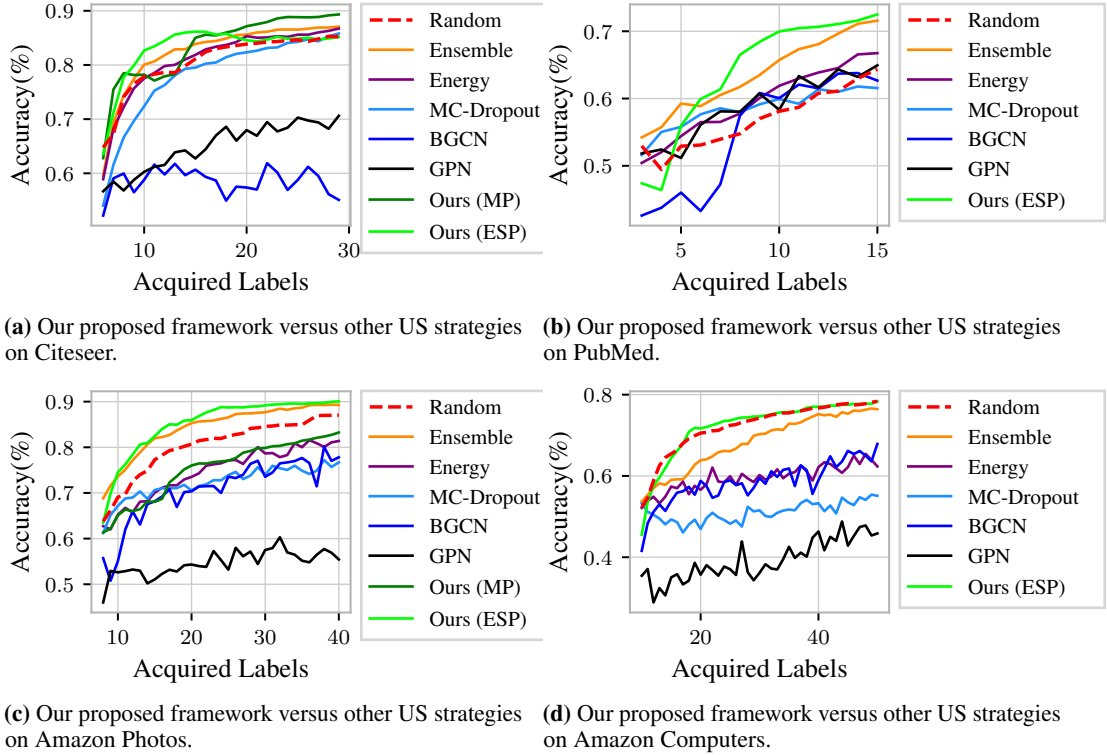
**(a)** AL using ground-truth uncertainty on a CSBM with 100 nodes and 7 classes. **(b)** AL using ground-truth uncertainty on a CSBM with 1000 nodes and 4 classes.



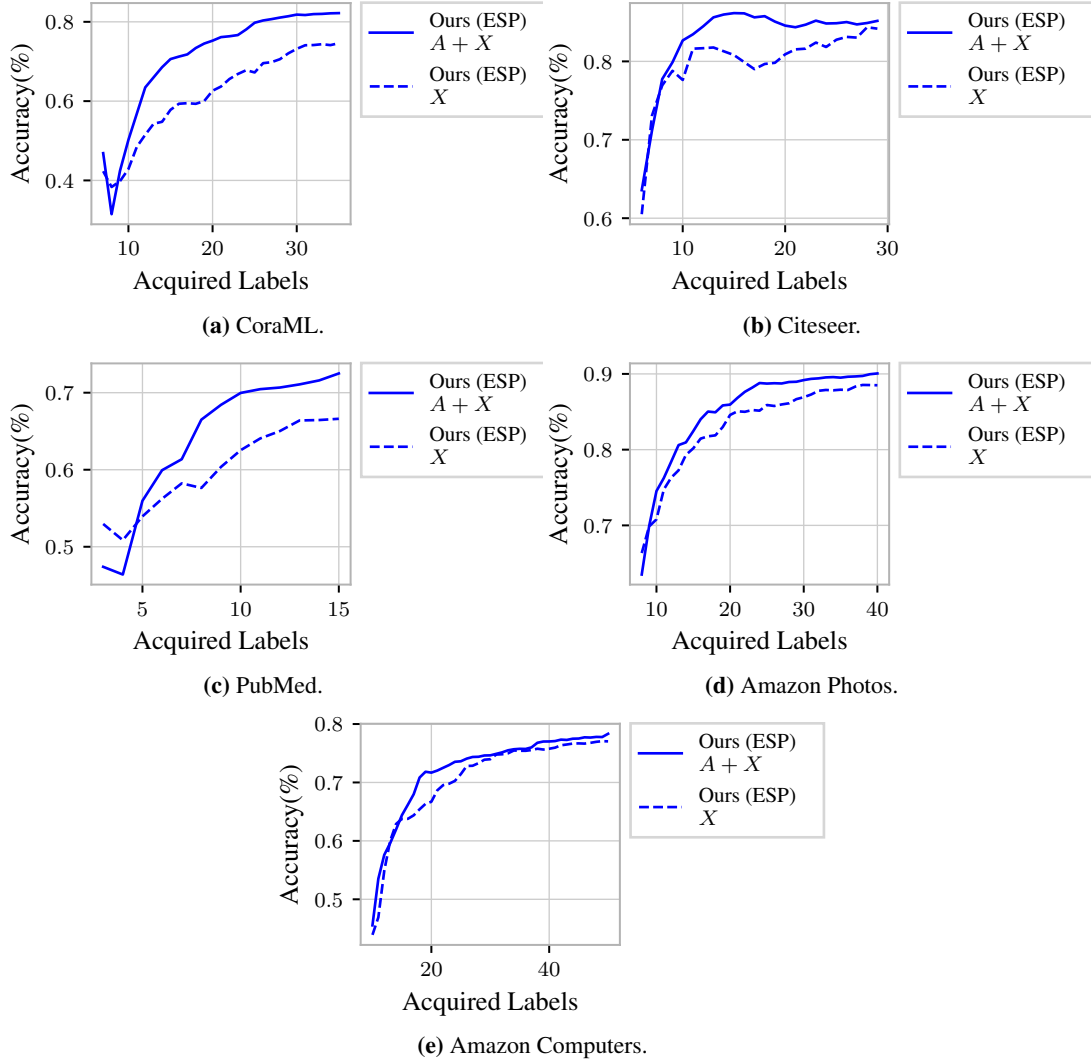
**(c)** AL using contemporary US strategies on a CSBM with 100 nodes and 7 classes. **(d)** AL using contemporary US strategies on a CSBM with 1000 nodes and 4 classes.



**Figure 8:** Evaluating the absolute improvement of Uncertainty Sampling using epistemic uncertainty over random acquisition for different structure and feature SNRs.



**Figure 12:** Our proposed uncertainty disentanglement framework applied to an SGC classifier using the MP or ESP approximation.



**Figure 13:** US using our proposed approximation (ESP) utilizing the features ( $A + X$ ) and network effects versus only using features ( $X$ ).