

# A Scalable and Adaptive System to Infer the Industry Sectors of Companies: Prompt + Model Tuning of Generative Language Models

Lele Cao\*, Vilhelm von Ehrenheim, Astrid Berghult, Cecilia Henje, Richard Anselmo Stahl, Joar Wandborg, Sebastian Stan, Armin Catovic, Erik Ferm and Hannes Ingelhag  
Motherbrain, EQT Group, Stockholm, Sweden

{lele.cao, vilhelm.vonehrenheim, astrid.berghult, cecilia.henje, richard.stahl}@eqtpartners.com  
{joar.wandborg, sebastian.stan, armin.catovic, erik.ferm, hannes.ingelhag}@eqtpartners.com

## Abstract

The Private Equity (PE) firms operate investment funds by acquiring and managing companies to achieve a high return upon selling. Many PE funds are thematic, meaning investment professionals aim to identify trends by covering as many industry sectors as possible, and picking promising companies within these sectors. So, inferring sectors for companies is critical to the success of thematic PE funds. In this work, we standardize the sector framework and discuss the typical challenges; we then introduce our sector inference system addressing these challenges. Specifically, our system is built on a medium-sized generative language model, finetuned with a prompt + model tuning procedure. The deployed model demonstrates a superior performance than the common baselines. The system has been serving many PE professionals for over a year, showing great scalability to data volume and adaptability to any change in sector framework and/or annotation.

## 1 Introduction

Private Equity (PE), as a fast-growing branch of the investment industry, operates investment funds on behalf of institutional and accredited investors by acquiring and managing companies before selling them to achieve high, risk adjusted returns. The common PE investment strategies, according to (Block et al., 2019), include venture capital, growth capital, and leveraged buyouts. The majority of PE funds strive to be “thematic” (Bérubé et al., 2014), aiming to identify macro-level trends by covering a variety of relevant sectors and picking promising companies within these sectors. In order to do that, any company should be put into a sector that best describes its main business activity. The sectors are often defined hierarchically (cf. the sector framework in Section 2), where the sectors higher up in the hierarchy tend to have a broader scope (hence

\*Corresponding author.

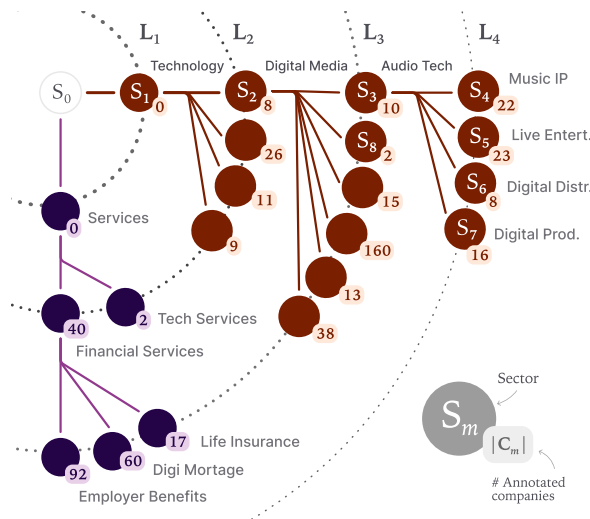


Figure 1: A PE sector framework defined as a tree with a depth  $L=4$ . Each non-root node represents a sector (i.e.,  $s_1 \sim s_M$ ) that is numbered in a depth-first order. The integer attached to the  $m$ -th sector/node indicates the number of companies  $|C_m|$  annotated for  $s_m$ .

usually fewer in number) and be more stable, while the ones lower down (a.k.a. “industries”) are more fine-grained and prone to change. A well-defined sector framework enables investment professionals to conduct a deeper analysis of the economy within each individual sector.

There are currently hundreds of millions of companies worldwide, and thousands of new companies are founded daily. Realistically, human professionals can only evaluate a limited number of companies to determine their belonging sectors. In order to significantly increase the coverage of sector mapping, practitioners have begun resorting to predictive systems to infer the belonging sectors of companies. Due to the reasons discussed in Section 2, there has not been any effective system that is generic enough to drive the wide adoption in PE operations. In this paper, we standardize the sector framework and discuss the typical challenges; we then introduce our sector inference system ad-

addressing these challenges. Our system has been successfully serving hundreds of PE professionals for over a year. The highlight is three fold:

- We propose to co-tune the PLM (pretrained language model) starting from a later stage of prompt tuning, attempting to leverage the capability of medium-sized PLMs to an extreme using scarce annotation.
- We implement an autonomous system, which effectively handles the dynamic sector framework, evolving annotation, data imbalance, noisy features, and high inference volume.
- We experimentally show the superior performance of our approach in comparison to the common baselines, and justify many design choices such as model paradigm and size.

## 2 The Problem and Challenges

Let  $c_n$  denote the  $n$ -th company ( $n = 1, 2, \dots, N$ ) in the scope of a PE firm; the total number of companies  $N$  usually reaches the order of millions. Most of the time, PE professionals maintain a hierarchical sector framework containing  $M$  different sectors represented as nodes ( $s_1 \sim s_M$ ) in a tree with  $L$  layers, as illustrated in Figure 1. In practice, the value of  $L$  is mostly less than 4, and the total number of sectors (i.e.,  $M$ ) tracked by a large PE firm may reach up to a few hundred. The problem is **how to assign each company  $c_n$  to the most relevant sector  $s_m$** . Solving such a problem requires addressing several **challenges** (abbreviated as **Chall.**) that will be discussed below.

**Chall.1:** scarce, imbalanced and evolving annotation. One might notice there are some public datasets such as G2 and Pitchbook<sup>1</sup> that contain sector annotations, i.e.,  $c_n \rightarrow s_m$ . In reality, they can not be directly used to train the sector inferring model, which is the consequence of two main facts: (1) PE firms almost always maintain their own version of sector framework that are drastically different from the ones from public datasets. (2) PE funds may annotate companies differently; for example, Klarna<sup>2</sup> might fall into any sector of *payment method*, *digital bank* and *financial service* depending on the preference of investment professionals or the fund specifications. To that end, we allow professionals to select a sector for

any company via the investment platform developed in-house. Formally, we use  $C_m$  to denote the set of companies annotated for sector  $s_m$ , and the total number of companies in  $C_m$  is  $|C_m|$ ; taking node  $(S_3)$  in Figure 1 for example, its subscript  $[10]$  contains the value of  $|C_3|$ , i.e.,  $|C_3| = 10$ . In reality, the sector annotation is scarce (an intrinsic limitation of manual annotation), imbalanced ( $|C_m|$  can vary greatly among sectors) and ever-evolving (the mapping  $c_n \rightarrow s_m$  may change frequently).

**Chall.2:** dynamic sector framework with varying granularity. Due to shifting market trends, the sector framework is rarely fixed for extended periods of time. Instead, the sector framework is really a dynamic one, where one of the three changes<sup>3</sup> can occur: adding new layers, adding new nodes, and removing nodes. Another observation is that PE professionals will pick concepts they think are important and define them as sectors, leading to sectors with varying granularity even on the same layer. For instance, a sector could be anything from a new technology (e.g., block chain), an environmental concern (e.g., water shortage), to an emerging market demand (e.g., Coronavirus test).

**Chall.3:** availability and quality of features. Intuitively, the most informative feature is probably the textual description about a company, which can be gathered from various data sources such as Pitchbook and Crunchbase<sup>4</sup>. Given an example description “We develop security analytical tools to identify web-app vulnerabilities. Contact us for a demo of our award-winning product”, one could guess a “cyber security” sector just by reading the first sentence, yet many texts look more like the second sentence, which severely lacks context. Moreover, a significant number of companies simply do not have textual descriptions available from popular data sources<sup>5</sup>.

**Chall.4:** high inference frequency and volume. As soon as the textual feature of a company is changed, we need to re-infer its sector. Besides, any change around the sector framework or company annotation may trigger model update, which

<sup>3</sup>Note that changing the definition of an existing sector is achieved by altering the associated company annotations; and merging/splitting existing sector(s) can be done via combining operations of adding and removing sector(s). Currently, only a system superuser can modify the sector framework through backend configuration files. However, our future plans involve facilitating this process via a web-based graphical user interface (GUI) integrated to EQT’s Motherbrain platform – <https://eqtgroup.com/motherbrain>.

<sup>4</sup><https://www.crunchbase.com>

<sup>5</sup>Cao et al. (2022) present a summary of data sources.

<sup>1</sup><https://www.g2.com> and <https://pitchbook.com>

<sup>2</sup><https://www.klarna.com>




Figure 2: Three paradigms of generative NLP models: prompting  $P(s|c; \theta_2)$ , model tuning  $P(s|c; \theta_1, \theta_2)$ , and prompt tuning  $P(s|c; \theta_2, \theta_3)$ .

requires a re-inference for all  $N$  companies. With the current data volume in our data warehouse, the daily amount of re-inference can easily exceed 100 million, which may grow into a bottleneck.

### 3 The Core Model

Inferring the industry sector of companies can be naturally addressed by a supervised NLP approach, where we **input** the textual description of a company (denoted as  $c$ ), and **output** a sector  $s$  based on a  $\theta$ -parameterized model  $P(s|c; \theta)$ ; note that we omit the subscripts  $n$  and  $m$  hereafter for the sake of simplicity. To find the optimal  $\theta$ , we use the annotated mappings  $c \rightarrow s$  to fit this conditional probability. The prediction target  $s$  can be either raw text (e.g., “cyber security”) or the encoded  $M$ -dim one hot vector, where the former is a **generative** approach and the latter is **discriminative**. It is crucial to highlight that generative methods offer two primary advantages over discriminative techniques (such as supervised classification): (1) generative models are capable of predicting sectors beyond those predefined, and (2) since these models output natural words, they can more effectively harness pre-learned knowledge in LM, thus avoiding overfitting on smaller training datasets.

Nowadays, generative approaches dominate the domains of computer vision (e.g., Stable Diffusion by Rombach et al., 2022) and NLP (e.g., GPT-3 (Brown et al., 2020) and GPT-4 (OpenAI, 2023)). Particularly, the language model (LM) is often pre-trained following a generative approach, such as predicting the masked words. To address **Chall.1&2**, we need to exploit the capability of a pretrained LM (PLM). We start with designing a template for samples:

[NAME], concerns [TAGS], is [c]. Sector: [s].

For a certain company, [NAME] is its legal name, [TAGS] is the concatenated tags/keywords<sup>6</sup> that are

<sup>6</sup>Many data sources, such as Pitbook and Crunchbase, have some keywords tagged for each company.

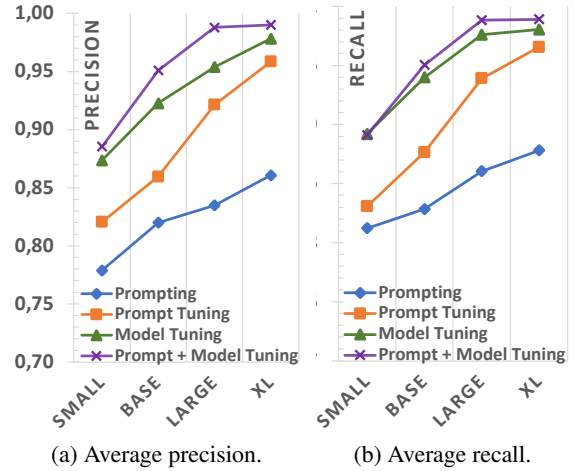


Figure 3: Performance comparison over different model paradigms (legend) and sizes (x-axis): the validation (a) precision and (b) recall averaged over 84 sectors.

added to address **Chall.3**. For example, the filled input for company “Klarna” may look like

*Klarna Bank AB, concerns buy-now-pay-later and shopping, is an online payment platform designed to facilitate cashless payments. Sector: [s].*

As the prediction target, [s] remains unreplaced, thus it is an unanswered sample. Model optimization essentially attempts to make the predicted [s] closer to the annotated sector text, and in this example  $s = \text{“financial service”}$ . We use  $c'$  to denote the filled sample for company  $c$ .

#### 3.1 Prompt and Model Tuning

Despite minor differences, the generative NLP models largely adhere to one of three paradigms: prompting, model tuning, or prompt tuning. Seen from Figure 2, **prompting** (Liu et al., 2023) freezes the PLM weights  $\theta_1$  while learning a mapping function (parameterized with  $\theta_2$ ) to transform the raw PLM output into the sector space. **Model tuning** allows finetuning  $\theta_1$ , which is the de facto way of leveraging large PLM for downstream tasks (Li and Liang, 2021). **Prompt tuning** prepends some soft prompts, which are essentially learnable virtual tokens, into the input sequence  $c'$  and only trains them (corresponding to parameter  $\theta_3$ ) while keeping  $\theta_1$  fixed (Su et al., 2022). PLM can have billions of parameters making model tuning paradigm expensive, while the prompt-based approach (Liu et al., 2023) has only thousands of tunable parameters (Lester et al., 2021).

Following the generative “text-to-text” T5 PLM (Raffel et al., 2020), we compared the performance of these paradigms towards the PE sector inferring

---

**Algorithm 1: Prompt + model tuning**

---

**Input:** Sector annotations in the form of  $c \rightarrow s$ , a generative NLP model  $P(s|c; \theta_1, \theta_2, \theta_3)$ , PLM freezing steps  $t'$ , learning rates  $\epsilon_1$  and  $\epsilon_2$

**Output:** The optimal parameters  $\theta_1^*$ ,  $\theta_2^*$  and  $\theta_3^*$

- 1 Initialize  $\theta_1$  by loading the pretrained T5 model;
  - 2 Initialize  $\theta_2$  and  $\theta_3$  randomly;
  - 3 **for** ( $t = 1; t \leq T; t++$ ) **do**
  - 4     Sample a mini-batch from the annotations;
  - 5     Transform each  $c$  into a filled template  $c'$ ;
  - 6     Forward propagate  $c'$  to obtain the prediction  $\hat{s}$ ;
  - 7     Calculate the T5 cross entropy loss  $\mathcal{L}(\hat{s}, s)$ ;
  - 8     **if**  $t \leq t'$  **then**
  - 9          $\epsilon = \epsilon_1$ ;
  - 10    **else**
  - 11          $\epsilon = \epsilon_2$  and  $\theta_1 := \theta_1 - \epsilon \frac{\partial \mathcal{L}(\hat{s}, s)}{\partial \theta_1}$ ;
  - 12          $\theta_2 := \theta_2 - \epsilon \frac{\partial \mathcal{L}(\hat{s}, s)}{\partial \theta_2}$  and  $\theta_3 := \theta_3 - \epsilon \frac{\partial \mathcal{L}(\hat{s}, s)}{\partial \theta_3}$ ;
  - 13  $\theta_1^* = \theta_1, \theta_2^* = \theta_2$  and  $\theta_3^* = \theta_3$ ;
  - 14 **return**  $\theta_1^*, \theta_2^*$  and  $\theta_3^*$ ;
- 

task. Figure 3 shows the average precision and recall in relation to different paradigms and model sizes (Small, Base, Large and XL)<sup>7</sup>. We observe that model tuning of T5 achieves stronger performance than prompting and prompt tuning. Prompt tuning catches up with model tuning as model size increases, which coincide the conclusion drawn by Lester et al. (2021). Intuitively, the label scarcity (Chall.1) and varying granularity of sector framework (Chall.2) could be better addressed by prompt tuning, since it is supposed to keep the learned knowledge in PLM untouched; meanwhile, we also want to replicate the superior performance of model tuning when using a smaller model. To that end, we propose a model  $P(s|c; \theta_1, \theta_2, \theta_3)$  that carries out  $t'$  steps of prompt tuning (only optimize  $\theta_2$  and  $\theta_3$ ) before jointly tune the PLM weights  $\theta_1$ , as presented in Algorithm 1. Seen from Figure 3, this ‘‘Prompt + Model Tuning’’ approach outperforms all compared methods by a large margin, which is the case even when the PLM size is relatively small.

According to Figure 3, the performance of our approach increase with the size of PLM and plateau (>98%) when reaching a ‘‘Large’’ size. Hence, we initialize our model with the T5-Large PLM and train for  $T = 1 \times 10^6$  steps with a mini-batch size of 50. The prompt tuning phase is trained for  $t' = 3 \times 10^3$  steps with a learning rate of  $\epsilon_1 = 0.1$ , where the first  $1 \times 10^3$  steps utilize a linear learning

<sup>7</sup>We did not manage to experiment the XXL T5 model due to our restriction of computing and human resources. The T5 PLMs can be found in <https://huggingface.co/google>.




Figure 4: Demonstration of annotation (label) attribution process using  $s_3$  sub-tree from Figure 1 as an example. Darker colored nodes are eligible for modeling.

rate warm-up (Goyal et al., 2017). Afterwards, the joint prompt and model tuning begins with a warm-up of  $1.5 \times 10^3$  steps until reaching a learning rate of  $\epsilon_2 = 5 \times 10^{-3}$ . Checkpoints are selected via early stopping with respect to the validation accuracy. All these hyper-parameters are determined by an empirical grid search, and the implementation is built upon OpenPrompt (Ding et al., 2022).

### 3.2 Annotation attribution

We empirically regulate that only the sectors with at least 20 annotated companies can be included in the modeling, which implies that some sectors, such as  $s_3$ ,  $s_6$  and  $s_7$  in Figure 1, are not eligible directly. Since sector annotation is scarce (Chall.1), we try to utilize every annotation to predict as much sectors as possible. Subsequently, we run a depth-first (bottom-up) annotation attribution algorithm to collect the eligible sectors  $s_m$  and their annotated set of companies  $C_m$ . Figure 4 demonstrates this procedure in three steps assuming the annotation attribution algorithm is currently processing the  $s_3$  sub-tree in Figure 1. Initially, only the child sectors  $s_4$  and  $s_5$  are eligible (cf.  $(S_4)_{22}$  and  $(S_5)_{23}$  in Figure 4) because they have more than 20 annotated companies. When it comes to  $s_6$  and  $s_7$ , they have insufficient annotations, thus are not eligible. However, their annotations will move up and contribute to the parent sector  $s_3$ , enabling  $s_3$  to be included in the training dataset due to  $|C_3| = 10 + 8 + 16 = 34 > 20$ .

Despite our best-effort annotation attribution procedure, it is possible that some sectors may still be excluded from training. However, in practice, the trained generative model is capable of producing sector names that are not within the eligible sector set. We believe this occurrence represents scenarios where the sectors are not covered by the labels, yet they are still significant in terms of their inherent business implications. This feature is particularly desirable as it facilitates better understanding and refinement of the sector framework.



**Algorithm 2:** Sample balancing via EDA

---

**Input:** The eligible sectors  $s_1, \dots, s_M$  and their corresponding company sets  $C_1, \dots, C_M$

**Output:** The balanced company sets  $C'_1, \dots, C'_M$

- 1 Initialize:  $C'_1 = C_1, \dots, C'_M = C_M$ ;
- 2 Calculate  $\zeta = 2 \cdot \max\{|C_1|, \dots, |C_M|\}$ ;
- 3 **for** ( $m = 1; m \leq M; m++$ ) **do**
- 4     **for each**  $c$  **in**  $C_m$  **do**
- 5         Augment  $c$  for  $\lfloor \zeta / |C_m| \rfloor - 1$  times with EDA (Wei and Zou, 2019), producing set  $c'$ ;
- 6          $C'_m = C'_m \cup c'$ ;
- 7 **return**  $C'_1, \dots, C'_M$ ;

---

### 3.3 Sample balancing via augmentation

As a part of **Chall.1** discussed in Section 2, the value of  $|C_m|$  can vary from merely 20 all the way to a few hundred. Thus, the aforementioned annotation attribution will produce a heavily imbalanced training dataset. The overall idea is augmenting the samples for minority sectors to achieve inter-sector balance. There is a whole spectrum of text augmentation methods: from rule-based to model-based techniques (Feng et al., 2021), from which we adopt the EDA (easy data augmentation) approach (Wei and Zou, 2019) because of its simplicity and universality. For individual [NAME], [TAGS] and [c] from our sample template, we perform synonym replacement, insertion, swapping and deletion at random choice with random intensity. Algorithm 2 has the details of the entire balancing procedure.

### 3.4 Performance analysis

As of December 2022, there are 84 eligible sectors after the annotation attribution procedure as introduced in Section 3.2. We collect all samples manually annotated under one of these 84 sectors, thereby creating a dataset that exhibits imbalance in terms of the number of samples annotated for each sector. The dataset is then balanced via the augmentation procedure introduced in Section 3.3. This results in a final dataset containing 7,260 samples, where each sector has  $\sim 86$  annotated samples in average. We reserve 15% of the dataset for validation and report the accuracy of different baselines in Table 1. Our approach (i.e., “Prompt + Model Tuning”) manages to achieve an accuracy of over 80% on the validation set. In contrast, its discriminative counterpart (cf. Section 5), which employs an  $M$ -way classification output head, achieves only 70% accuracy (largely on par with prompt tuning), likely due to the scarcity of labels.

Then we dig down to understand the error con-

Model	Accuracy (%)
$M$ -Way Classification	70.02
Prompting	64.63
Prompt Tuning	70.91
Model Tuning	76.44
Prompt + Model Tuning (Ours)	<b>80.25</b>

Table 1: Performance comparison of various baselines, all employing “T5 Large” as the PLM. The reported accuracies have been obtained (in December 2022) using the same random seed for consistency. The highest performing result is highlighted in bold.

L2	digital media	<b>12</b>	<b>0</b>	<b>0</b>	<b>0</b>	<b>0</b>	<b>2</b>	<b>1</b>
	deep tech.	<b>0</b>	<b>11</b>	<b>0</b>	<b>3</b>	<b>0</b>	<b>1</b>	<b>1</b>
L3	game	<b>1</b>	<b>0</b>	<b>14</b>	<b>0</b>	<b>0</b>	<b>0</b>	<b>1</b>
	cyber security	<b>0</b>	<b>0</b>	<b>0</b>	<b>15</b>	<b>0</b>	<b>0</b>	<b>0</b>
	market place	<b>0</b>	<b>0</b>	<b>0</b>	<b>0</b>	<b>13</b>	<b>0</b>	<b>2</b>
	horizontal software	<b>0</b>	<b>0</b>	<b>0</b>	<b>1</b>	<b>1</b>	<b>11</b>	<b>2</b>
	vertical software	<b>2</b>	<b>0</b>	<b>1</b>	<b>0</b>	<b>0</b>	<b>3</b>	<b>10</b>
	SECTOR →	digi. media	deep tech.	game	cyber sec.	market place	horiz. soft.	vert. soft.
	LAYER →	L2		L3				

Figure 5: The confusion matrix for seven sectors picked from the 84 (as of December 2022) predicted sectors.

tribution from each sectors and find that sectors on low levels (e.g., L3 and L4 in Figure 1) have an accuracy of over 90% except two L3 sectors named *horizontal software* and *vertical software*, as reflected in Figure 5. A horizontal software company caters to a wide and broad ranging market of consumers, and a vertical one provides a solution for a particular line of business or industry. Because of the way they are defined, many businesses in horizontal/vertical software sector might fit in other sectors as well. For example, a company providing bot-based customer service could be part of the horizontal market of any customer support scenario, while also targeting vertically to game publishers. The complete list of predicted sectors is considered to be sensitive proprietary information and therefore we only show the confusion matrix for seven sectors (two from L2 and five from L3) in Figure 5. Since L3 sectors are more fine-grained requiring less (than L2) annotations, a generally better inference performance is observed for L3 than L2, which encourages us to run a bottom-up annotation




Figure 6: The full system diagram with three scenarios (1~3) controlled by a rule-based inspection operator.

attribution (cf. Section 3.2) to prioritize lower-level sectors.

## 4 The Full System

The evolving annotation (**Chall.1**) and dynamic sector framework (**Chall.2**) both demand constant model iteration; and any model update would require a full-scale re-inference. As a result, triggering model iteration upon any change in annotation or sector framework will be computationally expensive and hard to scale. Instead, We trigger a rule-based inspection only once every day (e.g., at about 02:00 as exemplified in Figure 6) through Airflow<sup>8</sup>. The inspection rules lead to three scenarios:

- (1) **Finetune on T5 PLM** when the sector framework is changed or the annotation for any existing sector has evolved significantly<sup>9</sup>; it takes about 7 hours on  $2 \times$  Nvidia P100 GPU.
- (2) **Finetune on the latest sector model** when the sector annotation only changed marginally<sup>9</sup>; but the first scenario will be enforced after 90 days since its last execution.
- (3) **Skip finetune** otherwise and run incremental inference introduced in Section 4.2.

The second scenario takes less than 1/7 of the effort of the first scenario. We continue to present the key ingredients of finetune and inference.

### 4.1 Finetune

Figure 7 shows the finetune pipeline which is encapsulated in a docker<sup>10</sup> image run by Google Kubernetes Engine (GKE)<sup>11</sup>. From our data warehouse managed by BigQuery<sup>12</sup> (Melnik et al., 2010), the

<sup>8</sup>Apache Airflow: <https://airflow.apache.org>

<sup>9</sup>For sector  $s_m$  that originally has  $|C_m|$  annotated companies, the number of newly added/removed companies is  $\Delta_m$ , then  $\Delta_m/|C_m| \geq 0.75$  is regarded as **significant**, and  $0.75 > \Delta_m/|C_m| \geq 0.1$  is a **marginal** change.

<sup>10</sup><https://www.docker.com>

<sup>11</sup><https://cloud.google.com/kubernetes-engine>

<sup>12</sup><https://cloud.google.com/bigquery>




Figure 7: The finetune workflow. BQ: BigQuery, GCS: Google Cloud Storage, GKE: Google Kubernetes Engine.

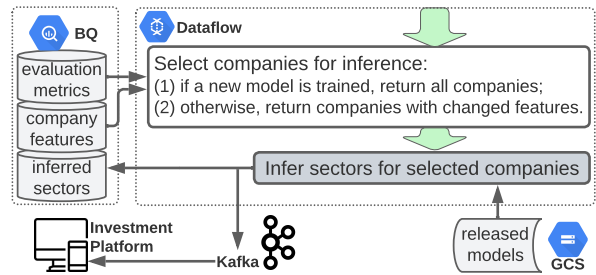


Figure 8: The (differentiated) inference workflow that starts with a rule-based company selection step.

annotation attribution (Section 3.2) collects all eligible sectors  $s_1, \dots, s_M$  together with their corresponding company sets  $C_1, \dots, C_M$ , which are balanced via augmentation (Algorithm 2). The balanced dataset is then split (with a ratio of 9:1) into training and validation sets that are used for prompt + model tuning following Algorithm 1.

The validation metrics (sector-wise precision and recall calculated from a confusion matrix like Figure 5) of the finetuned model will go through a QA (quality check) step to determine if this model is good enough to be automatically released in GCS (Google Cloud Storage)<sup>13</sup>. QA constitutes a series of assertions such as “*The precision of vertical software should be greater than 75%*” and so on. If any of these assertions fails, it will send an alarm to our data scientists via Slack<sup>14</sup> to request a manual interference to take appropriate actions.

### 4.2 Inference

The inference workflow starts with a selection step (cf. Figure 8) to determine a subset of companies that need re-inference. The selection step greatly reduce the daily inference load (by 95% at least after the system stabilizes), hence it tackles **Chall.4**.

<sup>13</sup><https://cloud.google.com/storage>

<sup>14</sup><https://slack.com>

Next, the latest trained model is loaded to infer the sectors for the selected companies, where two facts could relieve **Chall.4** further: (1) being able to use a medium-sized PLM, and (2) job parallelization by Dataflow<sup>15</sup>. The inferred sectors are stored in BigQuery and simultaneously published to Kafka<sup>16</sup> so that our investment platform can further leverage those predictions in many PE analytical use cases.

## 5 Related Work

As discussed in Section 3, the most relevant approach is  $M$ -way classification using either word or sentence level features as input. But the features are usually pre-learned with a fundamentally different setup and target (e.g., MLM: Masked Language Model, cf. Devlin et al., 2019), which makes it potentially difficult to continue finetuning towards a classification target (Gururangan et al., 2020). There is a recent trend of unifying all downstream tasks as a text generation problem (Lester et al., 2021), i.e., a generative NLP paradigm. In fact, using prompt has become the symbol of this paradigm. For example, we can potentially ask a GPT-3 Brown et al., 2020 or InstructGPT (Ouyang et al., 2022): “Klarna is a company that provide an cashless online payment platform. What is Klarna’s industry sector?” The likely answer would be something that generally make sense, yet will not be mapped directly towards the predefined sector framework. As a result, prompting (Liu et al., 2023) and prompt tuning (Su et al., 2022) emerge to fill this gap. However, Lester et al. (2021) discover that model tuning still prevails when the size of PLM is relatively small, which inspires us to jointly tune small PLM and prompt, as explained in Section 3.1.

## 6 Conclusion

In order to support thematic PE fund operations, we design and deploy a scalable and adaptive system to infer customized industry sectors for millions of companies. We empirically show that a generative NLP model is superior to its discriminative counterpart, leading to a solution of model + prompt tuning that guarantees superior performance even using scarce annotation and medium-sized PLM. The prompt template is designed to cope with noisy input textual features. To address the ever-changing sector framework and annotation, the system automatically triggers and determines

the most appropriate scenario by quantifying the change. Moreover, the system also incorporates best-effort annotation attribution, sample balancing, and incremental inference. Hundreds of PE professionals has benefited from this system for over a year. Last but not least, our solution can be directly generalized to many similar scenarios such as e-commerce product tagging.

## Acknowledgements

We are grateful to the support from the entire EQT organization. We also thank the constructive feedbacks from the reviewers of FinNLP@IJCAI 2023. This work is also reviewed by EQT’s compliance, communication and legal department prior to publication.

## References

- Vincent Bérubé, Sacha Ghai, and Jonathan Tétrault. 2014. *From indexes to insights: The rise of thematic investing*. *McKinsey Invest*, 1:51–56.
- Joern Block, Christian Fisch, Silvio Vismara, and René Andres. 2019. *Private equity investment criteria: An experimental conjoint analysis of venture capital, business angels, and family offices*. *Journal of corporate finance*, 58:329–352.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. *Language models are few-shot learners*. *Advances in neural information processing systems*, 33:1877–1901.
- Lele Cao, Vilhelm von Ehrenheim, Sebastian Krakowski, Xiaoxue Li, and Alexandra Lutz. 2022. *Using deep learning to find the next unicorn: A practical synthesis*. *arXiv preprint arXiv:2210.14195*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. *BERT: Pre-training of deep bidirectional transformers for language understanding*. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ning Ding, Shengding Hu, Weilin Zhao, Yulin Chen, Zhiyuan Liu, Haitao Zheng, and Maosong Sun. 2022. *Openprompt: An open-source framework for prompt-learning*. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 105–113.
- Steven Y. Feng, Varun Gangal, Jason Wei, Sarath Chandar, Soroush Vosoughi, Teruko Mitamura, and Edward Hovy. 2021. *A survey of data augmentation*

<sup>15</sup><https://cloud.google.com/dataflow>

<sup>16</sup><https://kafka.apache.org>

- approaches for NLP. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 968–988, Online. Association for Computational Linguistics.
- Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. 2017. [Accurate, large minibatch SGD: Training ImageNet in 1 hour](#). *arXiv preprint arXiv:1706.02677*.
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. [Don’t stop pretraining: Adapt language models to domains and tasks](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8342–8360, Online. Association for Computational Linguistics.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. [The power of scale for parameter-efficient prompt tuning](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3045–3059, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Xiang Lisa Li and Percy Liang. 2021. [Prefix-tuning: Optimizing continuous prompts for generation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4582–4597, Online. Association for Computational Linguistics.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2023. [Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing](#). *ACM Computing Surveys*, 55(9):1–35.
- Sergey Melnik, Andrey Gubarev, Jing Jing Long, Geoffrey Romer, Shiva Shivakumar, Matt Tolton, and Theo Vassilakis. 2010. [Dremel: interactive analysis of web-scale datasets](#). *Proceedings of the VLDB Endowment*, 3(1-2):330–339.
- OpenAI. 2023. [GPT-4 technical report](#). *arXiv preprint arXiv:2303.08774*.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. [Training language models to follow instructions with human feedback](#). *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. [High-resolution image synthesis with latent diffusion models](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695.
- Yusheng Su, Xiaozhi Wang, Yujia Qin, Chi-Min Chan, Yankai Lin, Huadong Wang, Kaiyue Wen, Zhiyuan Liu, Peng Li, Juanzi Li, Lei Hou, Maosong Sun, and Jie Zhou. 2022. [On transferability of prompt tuning for natural language processing](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3949–3969, Seattle, United States. Association for Computational Linguistics.
- Jason Wei and Kai Zou. 2019. [EDA: Easy data augmentation techniques for boosting performance on text classification tasks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6382–6388, Hong Kong, China. Association for Computational Linguistics.