

# ArchBERT: Bi-Modal Understanding of Neural Architectures and Natural Languages

Mohammad Akbari, Saeed Ranjbar Alvar, Behnam Kamranian,  
Amin Banitalebi-Dehkordi, Yong Zhang  
Huawei Technologies Canada Co., Ltd.  
{mohammad.akbari, saeed.ranjbar.alvar1, behnam.kamranian,  
amin.banitalebi, yong.zhang3}@huawei.com

## Abstract

Building multi-modal language models has been a trend in the recent years, where additional modalities such as image, video, speech, etc. are jointly learned along with natural languages (i.e., textual information). Despite the success of these multi-modal language models with different modalities, there is no existing solution for neural network architectures and natural languages. Providing neural architectural information as a new modality allows us to provide fast architecture-2-text and text-2-architecture retrieval/generation services on the cloud with a single inference. Such solution is valuable in terms of helping beginner and intermediate ML users to come up with better neural architectures or AutoML approaches with a simple text query. In this paper, we propose ArchBERT, a bi-modal model for joint learning and understanding of neural architectures and natural languages, which opens up new avenues for research in this area. We also introduce a pre-training strategy named Masked Architecture Modeling (MAM) for a more generalized joint learning. Moreover, we introduce and publicly release two new bi-modal datasets for training and validating our methods. The ArchBERT’s performance is verified through a set of numerical experiments on different downstream tasks such as architecture-oriented reasoning, question answering, and captioning (summarization). Datasets, codes, and demos are available as supplementary materials<sup>1</sup>.

## 1 Introduction

Existing machine learning models are mostly based on uni-modal learning, where a single modality is learned for the desired tasks. Example scenarios include image classification with image-only data; or language translation with text-only data (Raffel

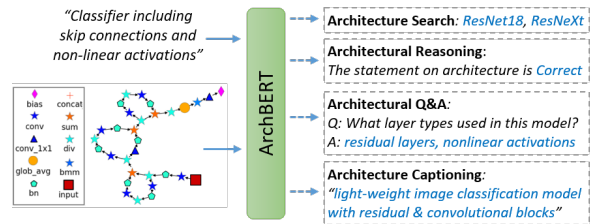


Figure 1: Bi-modal understanding of neural architectures and natural languages with sample applications.

et al., 2020; Akbari et al., 2022; Brown et al., 2020). Despite the success of existing uni-modal learning methods at traditional single-modal tasks, they are usually insufficient (Baltrušaitis et al., 2018) to model the complete aspects of human’s reasoning and understanding of the environment.

The alternative solution for this problem is to use multi-modal learning, where a model can jointly learn from multiple modalities such as text, image, or video to yield more abstract and generalized representations. As a result, a better understanding of various senses in information can be achieved and many new challenges that concern multi-modality can be handled. Such solution also enables the possibility of supplying a missing modality based on the observed ones. As an example, in text-based image generation, we aim to generate photo-realistic images which are semantically consistent with some given text description (Bao et al., 2022).

One of the most popular multi-modal solutions is multi-modal language models (LMs), where an extra modality (e.g., image or video) is jointly used and learned along with the natural languages (i.e., textual information). Some of the recent multi-modal LMs include ViLBERT for image+text (Lu et al., 2019), VideoBERT for video+text (Sun et al., 2019), CodeBERT for code+text (Feng et al., 2020), and also GPT-4 (OpenAI, 2023).

Although many multi-modal LMs with different modalities have been introduced so far, there is no existing solution for joint learning of neural

<sup>1</sup><https://developer.huaweicloud.com/develop/aigallery/notebook/detail?id=e6a924c7-735a-4e02-a25b-4416b77b6315>

network architectures and natural languages. Providing neural architectural information as a new modality allows us to perform many architecture-oriented tasks such as Architecture Search (AS), Architecture Reasoning (AR), Architectural Question Answering (AQA), and Architecture Captioning (AC) (Figure 1). The real-world applications of such solution include fast architecture-2-text and text-2-architecture retrieval/generation services on the cloud with a single inference. Such solution is valuable in terms of helping users to come up with better neural architectures or AutoML approaches with a simple text query especially for beginner and intermediate ML users. For instance, AC can be used for automatically generating descriptions or model card information on a model hub (i.e., machine learning models repository). Furthermore, AR is helpful when a model is uploaded to a repository or cloud along with some textual description provided by the user, where the relevancy of the user’s description for the given model can be automatically verified. If not verified, alternative auto-generated descriptions by a architecture-2-text solution can be proposed to the user.

In this paper, we propose ArchBERT as a bi-modal solution for neural architecture and natural language understanding, where the semantics of both modalities and their relations can be jointly learned (Figure 1). To this end, we learn joint embeddings from the graph representations of architectures and their associated descriptions. Moreover, a pre-training strategy called Masked Architecture Modelling (MAM) for a more generalized and robust learning of architectures is proposed. We also introduce two new bi-modal datasets called TVHF and AutoNet for training and evaluating ArchBERT. To the best of our knowledge, ArchBERT is the first solution for joint learning of architecture-language modalities. In addition, ArchBERT can work with any natural languages and any type of neural network architectures designed for different machine learning tasks. The main contributions of this paper are as follows:

- A novel bi-modal model for joint learning of neural architectures and natural languages
- Two new bi-modal benchmark datasets for architecture-language learning and evaluation
- A new pre-training technique called MAM
- Introducing and benchmarking 6 architecture-language-related downstream applications

## 2 Related Works

Multi-modal models are used in many sub-fields in machine learning. For example, [Michelsanti et al. \(2021\)](#) and [Schoneveld et al. \(2021\)](#) introduced the audio-visual models trained on input acoustic speech signal and video frames of the speaker for speech enhancement, speech separation, and emotion recognition. Multi-modal models used in biomedical ([Venugopalan et al., 2021](#); [Vale-Silva and Rohr, 2021](#)), remote-sensing ([Hong et al., 2020](#); [Maimaitijiang et al., 2020](#)), and autonomous driving ([Xiao et al., 2020](#)) applications have also proven to provide more accurate prediction and detection than the unimodal models.

Among different types of multi-modal LMs in the literature, transformer-based ones have shown significant performance, especially for vision-and-language tasks like visual question answering, image captioning, and visual reasoning. In VisualBERT ([Li et al., 2019](#)), a stack of transformers is used to align the elements of text and image pairs. ViLBERT ([Lu et al., 2019](#)) extended BERT to a multi-modal double-stream model based on co-attentional transformer layers. In LXMERT ([Tan and Bansal, 2019](#)), three encoders including language, object relation, and cross modality encoders are used. A single-stream vision-language model was introduced in VL-BEIT ([Bao et al., 2022](#)), where unpaired and paired image-text modalities were used for pre-training.

Video is another modality that is used with language in multi-modal models. VideoBERT ([Sun et al., 2019](#)) is a single-stream video-language model, which learns a joint visual-linguistic representation from input video-text pairs. VIOLET ([Fu et al., 2021](#)) is another example that employs a video transformer to model the temporal dynamics of videos, and achieves SOTA results on video question answering and text-to-video retrieval. Programming language is also an emerging modality that has been used along with language. For example, CodeBERT ([Feng et al., 2020](#)) is a multi-stream model, which uses LMs in each stream, where the input code is regarded as a sequence of tokens. On the other hand, GraphCodeBERT ([Guo et al., 2021](#)) proposes a structure-aware pre-training technique to consider the inherent structure of the code by mapping it to a data flow graph.

There are several prior works that combine more than two modalities. In Multimodal Transformer (MulT) ([Tsai et al., 2019](#)), cross-modal attention

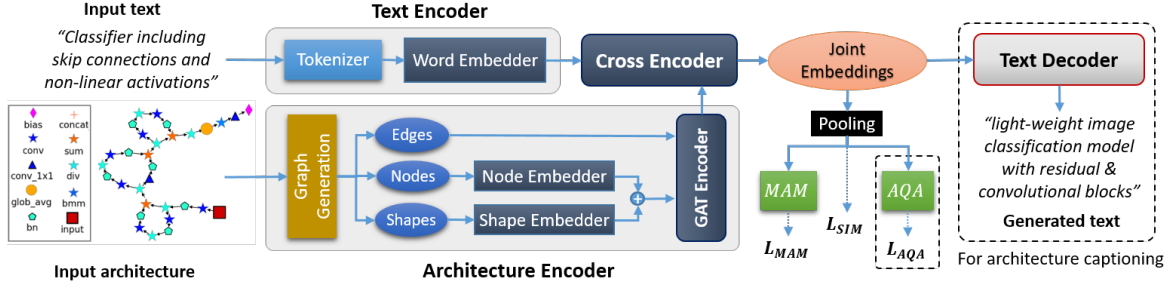


Figure 2: Overall framework of ArchBERT.

modules are added to the transformers to learn representations from unaligned multi-modal streams, including the language, the facial gestures, and the acoustic behaviors. VATT (Akbari et al., 2021) also used video, audio, and text transformers along with a self-supervised learning strategy to obtain multi-modal representations from unlabeled data.

It is worth mentioning that ChatGPT (OpenAI, 2022) can be used for information retrieval, question answering, and also summarization over the textual descriptions of well-known neural architectures such as AlexNet (Krizhevsky et al., 2017) or Faster-RCNN (Ren et al., 2015). However, unlike ArchBERT, it does not have a bi-modal understanding of both neural architectures (i.e., graphs) and natural languages, especially for newly proposed architectures and models.

### 3 Proposed Method: ArchBERT

The overall ArchBERT framework is shown in Figure 2. The major components of ArchBERT include a text encoder, an architecture encoder, a cross encoder, and a pooling module.

First, the input text represented by a sequence of  $n$  words  $W = \{w_i | i \in [1, n]\}$  is tokenized to a sequence of  $n$  tokens  $T = \{t_i | i \in [1, n]\}$ . Then, the text encoder  $E_t$  is utilized to map them to some word/token embeddings denoted by  $M_t \in \mathbb{R}^{(n \times d)}$  with the embedding size of  $d$ :  $M_t = E_t(T)$ .

On the other hand, the architecture encoder is responsible for encoding the input neural architecture. In this procedure, the computational graph of the input architecture is first extracted and represented with a directed acyclic graph  $G = \{V, A, S\}$  where  $V = \{v_i | i \in [1, m]\}$  denotes a sequence of  $m$  nodes representing the operations and layers (e.g., convolutions, fully-connected layers, summations, etc.) and  $A \in \{0, 1\}^{m \times m}$  denotes a binary adjacency matrix describing the edges and the connectivity between the nodes. In addition to the nodes and edges, we also extract the shape of the param-

eters associated with each node (i.e., input/output channel dimensions and kernel sizes), denoted by  $S = \{(s_i \in \mathbb{N}^4) | i \in [1, m]\}$ .

The nodes and the shapes are separately encoded using the node and shape embedders  $E_v$  and  $E_s$ , respectively. The adjacency matrix along with the summation of the resulting nodes and shapes embeddings are then given to a Graph Attention Network (GAT) (Veličković et al., 2018) for computing the final architecture (graph) embeddings denoted by  $M_g \in \mathbb{R}^{(m \times d)}$  with the embedding size of  $d$ :

$$M_g = GAT(E_v(V) + E_s(S), A) \quad (1)$$

In general, GAT is designed to operate on graph-structured data in which a set of graph features (node+shape embeddings in our case) is transformed into higher-level features. Given the adjacency matrix, the GAT model also allows all nodes to attend over their neighborhoods' features based on a self-attention strategy.

For joint learning of textual and architectural embeddings and share learning signals between both modalities, a cross transformer encoder,  $E_c$ , is used to process both embeddings in parallel. These embeddings are then average-pooled to fixed-size 1D representations  $J_t \in \mathbb{R}^{(1 \times d)}$  and  $J_g \in \mathbb{R}^{(1 \times d)}$ :

$$\{J_t, J_g\} = E_c(\{M_t, M_g\}) \quad (2)$$

As in S-BERT (Reimers and Gurevych, 2019), we use the cosine similarity loss as a regression objective function to learn the similarity/dissimilarity between architectures and language embeddings. First, the cosine similarity between  $J_t$  and  $J_g$  are computed. Given a target soft score  $y \in [0, 1]$  (i.e., 0: dissimilar, 1: similar), the following mean squared-error (MSE) loss is then employed:

$$L_{SIM} = \|y - \frac{J_t \cdot J_g}{\max(\|J_t\|_2, \|J_g\|_2, \epsilon)}\|_2, \quad (3)$$

which minimizes the cosine distance between  $J_t$  and  $J_g$  pairs labeled as similar, while maximizes the distance for the dissimilar ones.

### 3.1 Masked Architecture Modeling (MAM)

In the literature, a well-known pre-training objective function called Masked Language Modeling (MLM) is widely used by BERT-based models for learning language representations (Devlin et al., 2019). Inspired by MLM, we introduce a new objective called Masked Architecture Modeling (MAM) to provide more generalized learning and understanding of the graph embeddings corresponding to the neural architectures by ArchBERT.

Inspired by BERT (Devlin et al., 2019), we randomly mask 15% of the nodes with a special mask token and re-produce the masked nodes under the condition of the known ones. The MAM objective function is then defined as:

$$L_{MAM} = -\mathbb{E}_{V_i \sim V} \log p(V_i | \hat{V}), \quad (4)$$

where  $\hat{V}$  is the masked version of  $V$ . In other words,  $\hat{V}$  includes the contextual unmasked tokens surrounding the masked token  $V_i$ . In practice, the corresponding probability distribution is obtained by the MAM head  $H_M$ . The MAM head defines the distribution by performing the softmax function on the logits  $F_m \in \mathbb{R}^{(m \times |\mathcal{E}|)}$  mapped from the graph embeddings  $J_g$  as follows:  $F_m = H_M(J_g)$ , where  $\mathcal{E}$  is the entire vocabulary of nodes (or nodes corpus) set. Given  $L_{SIM}$  and  $L_{MAM}$ , the following weighted loss is then used for optimizing and pre-training the ArchBERT model:

$$L = L_{SIM} + \alpha L_{MAM}. \quad (5)$$

### 3.2 Architectural Question Answering (AQA)

The pre-trained ArchBERT can be utilized for the AQA task that is defined as the procedure of answering natural language questions about neural architectures. In other words, we can enable the ArchBERT model to predict the answers to architecture-related questions when the architecture and the question are matched.

For this task, we can fine-tune ArchBERT as a fusion encoder to jointly encode the input neural architecture and the question. To this end, the question and the architecture are first encoded using the text and architecture encoders, respectively. Both embeddings are then cross-encoded and pooled in order to calculate the final joint embeddings  $J_t$  and  $J_g$ . The element-wise product is then computed to interactively catch similarity/dissimilarity and discrepancies between the embeddings. The resulting product is fed into AQA head for mapping to the logits  $F_q \in \mathbb{R}^{|\mathcal{A}|}$  corresponding to  $|\mathcal{A}|$  answers:

$$F_q = H_q(J_t, J_g) \quad (6)$$

As in (Anderson et al., 2018), the AQA in our work is formulated as a multi-label classification task, which assigns a soft target score to each answer based on its relevancy to  $|\mathcal{A}|$  answers. A binary cross-entropy loss (denoted by  $L_{AQA}$ ) on the target scores is then used as objective function.

### 3.3 Language Decoder

We can empower the pre-trained ArchBERT to learn from and then benefiting for neural architecture captioning (or summarization) task by attaching a transformer decoder (Lewis et al., 2020) to generate textual tokens one by one. In this regard, an auto-regressive decoding procedure is employed with the following loss function:

$$L_{DEC} = -\mathbb{E}_{T_i \sim T} \log p(T_i | T_{<i}, \hat{T}), \quad (7)$$

where  $\hat{T}$  is the masked version of the ground truth text  $T$ , and  $T_i$  is the  $i$ -th token to be predicted.  $T_{<i}$  denotes the set of all the tokens decoded before  $T_i$ . Similar to MAM, the probability distribution over the whole vocabulary is practically obtained by applying softmax on the decoded feature (or logits)  $F_d \in \mathbb{R}^{(m \times |\mathcal{C}|)}$  that is calculated by providing the graph embeddings  $J_g$  to the decoder:  $F_d = D_t(J_g)$ , where  $\mathcal{C}$  denotes the entire vocabulary set.

## 4 Datasets

For pre-training the ArchBERT model, a dataset of neural architectures labeled with some relevant descriptions is required. To the best of our knowledge, there is no such bi-modal dataset in the literature. In this paper, we introduce two datasets called TVHF and AutoNet for bi-modal learning of neural architectures and natural languages. The numerical and the statistical details of TVHF and AutoNet datasets are summarized in Table 1.

Note that all the labels and descriptions in the proposed datasets have been manually checked and refined by human. There may be some minor noise in the dataset (i.e., an inevitable nature of any dataset, especially the very first versions), but in overall, the datasets are of sufficient quality for our proof-of-concept experiments.

### 4.1 TVHF

In order to create this dataset, we collected 538 unique neural architectures from TorchVision (TV) (Marcel and Rodriguez, 2010) and HuggingFace (HF) (Wolf et al., 2019) frameworks. The descriptions relevant to the architectures were extracted

Table 1: Statistical details of TVHF and AutoNet datasets (\*: AQA,  $\mu$ : mean,  $\sigma$ : standard deviation,  $M$ : median).

Dataset	Split	#Samples	Architecture									Text					
			#Unique Archs	#Unique Nodes	#Nodes			#Edges			#Unique Tokens	#Tokens			Sequence Length		
					$\mu$	$\sigma$	$M$	$\mu$	$\sigma$	$M$		$\mu$	$\sigma$	$M$			
TVHF	Train	24069	538	50	1146.61	1162.38	705	1281	1302.90	753	3507	16.16	11.22	14	97.60	77.76	81
	Val	6018	538	50	1146.61	1162.38	705	1281	1302.90	753	2965	16.21	11.59	14	97.88	80.33	81
AutoNet	Train	103306	10000	28	371.50	312.61	266	401	322.99	241	769	43.81	8.62	45	333.67	74.80	345
	Val	10338	1000	28	384.48	343.31	266	419	368.20	293.5	652	43.92	8.66	45	334.01	74.92	345
AutoNet*	Train	350000	10000	28	373.33	313.90	270	404	325.45	297	86	10.78	1.89	11	62.76	12.48	62
	Val	35000	1000	28	358.3	301.98	261	390	324.31	285.5	86	10.79	1.89	11	62.76	12.45	62

from TV and HF frameworks as well as other on-line resources such as papers and web pages (with the vocabulary size  $|\mathcal{C}|=31,764$ ). To increase the dataset size, the descriptions were split into individual sentences each assigned to the related architecture, which provided a collection of 2,224 positive samples, i.e., pairs of architecture with their relevant descriptions (details in the appendix).

To assure the model learns both similarities and dissimilarities, we also generated negative samples by assigning irrelevant descriptions to the architectures (resulting in a total of 27,863 negative samples). We randomly split the dataset (in total 30,087 samples) into 80% for train and 20% for validation.

For fine-tuning and evaluating ArchBERT on Architecture Clone Detection (ACD), we establish another dataset including pairs of architectures manually hard-labeled with a dissimilarity/similarity score (0 or 1). To this end, all combinations of two architectures from TVHF were collected (in total 82.8K samples) and split into train/val sets (80% and 20%). Details are provided in the appendix.

## 4.2 AutoNet

As described before, TVHF includes realistic human-designed architectures, which are manually labeled with real descriptions. On the other hand, we introduce the AutoNet dataset, which includes automatically generated architectures and descriptions. AutoNet is basically the modified and extended version of DeepNet1M (Knyazev et al., 2021), which is a standardized benchmark and dataset of randomly generated architectures for the parameter prediction tasks.

In AutoNet, we extend the set of operations (layers) from 15 types (in DeepNet1M) to 85, which include most of the recent operations used in computer vision and natural language models. We followed the same procedure in DeepNet1M and ran-

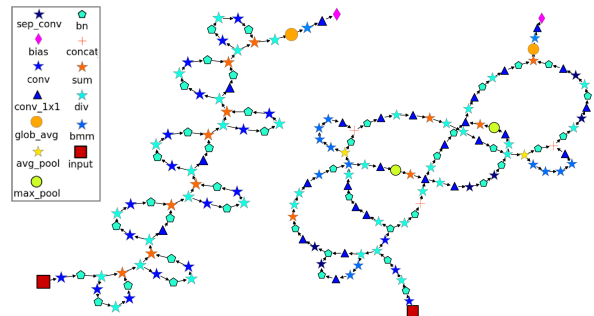


Figure 3: Sample graphs generated for ResNet18 (left) and a random architecture from AutoNet (right).

domly generated 10K and 1K architectures for train and validation sets, respectively.

For automatic generation of textual descriptions related to each architecture, we created an extensive set of sentence templates, which were filled based on the information extracted from the structure, modules, and existing layers of the corresponding architecture. The same process was applied for generating negative samples, but with the textual information of the non-existing modules and layers in the architecture. For each architecture, 10-11 textual descriptions were created, which resulted in 103,306 and 10,338 architecture and text pairs for the train and validation sets (with the vocabulary size  $|\mathcal{C}|=30,980$ ), respectively. The details of this procedure are given in the appendix.

### 4.2.1 AutoNet-AQA

For fine-tuning and evaluating ArchBERT on AQA, another dataset including triplets of architectures, questions, and answers is needed. As in AutoNet, a set of question/answer templates were used to automatically generate the questions and answers. The same procedure of generating neural architectures as in AutoNet was employed. 10K and 1K architectures were respectively created for the train and validation sets. For each architecture, 35 unique

questions were generated, and the answers were chosen from a list of  $|\mathcal{A}| = 51$  unique answers. In total, the train and validation sets respectively include 350K and 35K samples.

The visualization of two sample graphs generated for ResNet18 from TVHF and a random architecture from AutoNet is shown in Figure 3. More sample data along with the quality analysis of the datasets are given in the appendix.

## 5 Experimental Results

In this section, the performance of ArchBERT on the following downstream tasks is evaluated and numerically analyzed.

- **Architectural Reasoning (AR)**: it is the task of determining if a statement regarding an architecture is correct or not.
- **Architecture Clone Detection (ACD)**: it includes the process of checking if two architectures are semantically/structurally similar or not.
- **Architectural Question Answering (AQA)**: as given in Section 3, it is the process of providing an answer to a question over a given architecture.
- **Architecture Captioning (AC)**: it is the task of generating descriptions for a given architecture.

Since there is no related prior works, we compare our method with some uni-modal baselines for each of the above tasks. An ablation study over different components of ArchBERT is also presented.

In this work, we employ the BERT-Base model (with 12 heads) as our ArchBERT’s cross encoder. We pre-trained ArchBERT on both TVHF and AutoNet datasets with a batch size of 80, embedding size of  $d=768$ , and the Adam optimizer with learning rate of  $2e-5$  for 6 hours. The training on TVHF and AutoNet was respectively done for 20 and 10 epochs. Since there is a large scale difference between the  $L_{SIM}$  and  $L_{MAM}$  loss values in the weighted loss in Equation 5, where  $L_{MAM} \gg L_{SIM}$ , we set  $\alpha=5e-2$  to balance the total loss value (obtained experimentally). A batch size of 80 is used for all the tests with the pre-trained ArchBERT.

### 5.1 Uni-Modal Baselines

For the AR baseline, we compare the architecture name with an input statement, which is considered as "correct" if the architecture name appears in the statement, otherwise it is "incorrect". Note that unlike this baseline, ArchBERT does not need the architecture name to infer about the statements.

For the ACD uni-modal baseline (Figure 4-left), the architecture encoder is first used to separately map both input architectures, denoted by  $\{G^1, G^2\}$ , into the graph embeddings  $\{M_g^1, M_g^2\}$  (Equation 1). The cross encoder and pooling module are then applied to obtain the fixed-size joint representations  $\{J_g^1, J_g^2\}$  (Equation 2). The cosine similarity loss in Equation 3 is finally performed on  $\{J_g^1, J_g^2\}$  pairs along with a provided hard-label. For this baseline, we trained ArchBERT with architecture-only pairs (without text encoder) from TVHF-ACD train set.

For the AQA uni-modal baseline (Figure 4-middle), we train a text-only ArchBERT (without architecture encoder), where the context is obtained from the textual information and summary of the input architecture, e.g., layer names (i.e., using Pytorch model summary function). The extracted information is considered as the input context on which the question answering procedure is performed. The tokenized input question and context, denoted by  $\{T^q, T^c\}$ , are mapped into token embeddings  $\{M_t^q, M_t^c\}$ , which are then cross-encoded and average-pooled to obtain the joint embeddings  $\{J_t^q, J_t^c\}$  (Equation 2). As in Equation 6, the element-wise product of  $\{J_t^q, J_t^c\}$  is given to the AQA head to obtain the logits required for the binary cross-entropy loss described in Section 3.2.

For the AC uni-modal baseline (Figure 4-right), we trained ArchBERT (without text encoder) followed by the decoder from scratch (no bi-modal pre-training of ArchBERT). The detailed AC procedure is described in Section 3.3.

### 5.2 Architectural Reasoning (AR)

For this task, the input text and the architecture are given to ArchBERT to create the pooled embeddings. The cosine similarity score between these embeddings is then computed. If the score is greater than some threshold  $\tau$  (i.e., 0.5), the statement on the architecture is determined as “correct”, otherwise “incorrect”. We evaluate the performance of the pre-trained ArchBERT on this task over the TVHF validation set. As summarized in Table 2, an accuracy and F1 score of 96.13% and 71.86% were respectively achieved. F1 scores are reported to deal with the class imbalance.

As reported in Table 2, a F1 score of 55.93% is achieved by the AR baseline, which is about 16% lower than ArchBERT.

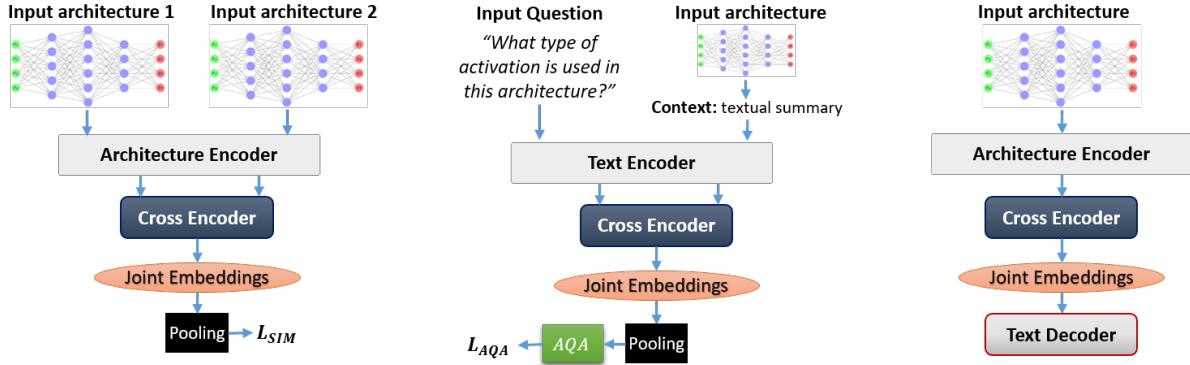


Figure 4: Uni-Modal Baselines (left: ACD, middle: AQA, right: AC).

### 5.3 Architecture Clone Detection (ACD)

To perform this task, both input architectures are given to ArchBERT’s architecture encoder followed by the cross-encoder and pooling module to obtain the pooled embeddings. The cosine similarity of the embeddings is then computed. If the similarity score is greater than a threshold (i.e., 0.5), the two architectures are considered similar, otherwise dissimilar.

We first evaluate the pre-trained ArchBERT’s performance on the TVHF-ACD validation set. Although the pre-trained model has not specifically learned to detect similar/dissimilar architectures, it still achieves a good accuracy of 86.20% and F1 score 60.10% (Table 2). However, by fine-tuning the pre-trained ArchBERT with TVHF-ACD train set, significantly improved accuracy and F1 score of 96.78% and 85.98% are achieved.

Two baselines including Jaccard similarity (Santesteban and Tejada-Cárcamo, 2015) and a uni-modal version of ArchBERT are used to compare with our bi-modal ArchBERT on ACD task. For Jaccard, the similarity of the architecture pairs is computed by taking the average ratio of intersection over union of the nodes and edges ( $V$  and  $A$ ). The pairs are considered as "similar" if the similarity score is greater than 0.5, otherwise "dissimilar". As shown in Table 2, the pre-trained and fine-tuned ArchBERT models respectively outperform this baseline with 14% and 40% higher F1 scores. The ACD uni-modal baseline also achieves F1 score of 84%, i.e., 2% lower than fine-tuned ArchBERT.

### 5.4 Architectural Question Answering (AQA)

For this, ArchBERT along with the attached AQA head (composed of a two layer MLP) is fine-tuned with the AutoNet-AQA dataset using a batch size of 140 over 10 epochs (for about 10 hours). We use

Table 2: The performance of ArchBERT and its components on different tasks and datasets (AR: Architectural Reasoning, ACD: Architecture Clone Detection, AQA: Architectural Question Answering, CR: Cross Encoder, MAM: Masked Architecture Modeling).

Task	Dataset	Model	Acc(%)	F1(%)
AR	TVHF	<b>ArchBERT</b>	<b>96.13</b>	<b>71.86</b>
		-w/o Shape	95.44	69.16
		-w/o Edge	95.52	68.98
		-w/o Edge+Shape	95.12	65.80
		-w/o MAM	95.18	64.27
		-w/o CR	94.42	57.03
		Baseline	89.03	55.93
ACD	TVHF	<b>ArchBERT</b>	<b>86.20</b>	<b>60.10</b>
		-w/o Shape	85.44	60.20
		-w/o Edge	76.70	47.96
		-w/o Edge+Shape	82.90	56.45
		-w/o MAM	78.80	49.59
		-w/o CR	69.89	42.35
		Jaccard	80.22	45.96
<b>ArchBERT-ft</b>	<b>96.78</b>	<b>85.98</b>		
Baseline (uni)	96.24	84.01		
AQA	AutoNet	<b>ArchBERT</b>	<b>72.73</b>	<b>73.51</b>
		-w/o MAM	66.08	66.16
		-w/o CR	60.32	63.33
		Baseline (uni)	55.82	61.84

the Adam optimizer with an initial learning rate of  $2e-5$ . At the inference time, we simply take a sigmoid over the AQA head’s logits (with the same batch size of 140). As given in Table 2, ArchBERT achieves an accuracy of 72.73% and F1 score of 73.51% over the AutoNet-AQA validation set.

For the AQA baseline, an F1 score of 61.84% was obtained on AutoNet-AQA, which is  $\approx 12\%$  lower than the proposed bi-modal ArchBERT.

### 5.5 Architecture Captioning (AC)

To analyze ArchBERT’s performance on AC, the pre-trained ArchBERT (without text encoder) attached with a language decoder is fine-tuned on both TVHF and AutoNet with a batch size of 30 for

Table 3: ArchBERT’s performance on Architecture Captioning (AC) (CR: Cross Encoder, MAM: Masked Architecture Modeling, R1: Rouge1-Fmeasure, R2: Rouge2-Fmeasure, RL: Rouge-Lsum-Fmeasure).

Dataset	Model	R1	R2	RL
TVHF	<b>ArchBERT</b>	<b>0.18</b>	0.05	<b>0.17</b>
	-w/o MAM	0.17	0.05	0.15
	Baseline (uni)	0.18	0.07	0.17
AutoNet	<b>ArchBERT</b>	<b>0.48</b>	<b>0.36</b>	<b>0.46</b>
	-w/o MAM	0.45	0.34	0.43
	Baseline (uni)	0.40	0.30	0.38

10 epochs. The fine-tuning process for TVHF and AutoNet respectively took about 0.5 and 6 hours. Adam optimizer with an initial learning rate of  $2e-5$  was used. For the language decoder, a single-layer transformer decoder (with 12 heads and hidden size of  $d=768$ ) followed by 2 linear layers is used.

At the inference, the beam search (with the size of 10) was employed to auto-regressively generate the output tokens, which were then decoded back to their corresponding words. The same batch size of 30 was used for the evaluation. The results over the TVHF and AutoNet validation sets are summarized in Table 3, where Rouge-Lsum-Fmeasure (RL) (Lin, 2004) scores of 0.17 and 0.46 were respectively achieved. Unlike AutoNet, TVHF dataset includes more complicated neural architectures along with high-level human-written textual descriptions, which makes the architecture captioning more challenging. As a result, lower performance is achieved.

The uni-modal AC baseline achieves an RL of 0.38 on AutoNet, which is 8% lower than the proposed bi-modal ArchBERT (i.e., pre-trained on both architectures and text, and fine-tuned for AC).

## 5.6 Architecture Search (AS)

ArchBERT is also applicable to Architecture Search (AS) downstream task. The task is to design a semantic search engine to receive a textual query from the user, search over a database of numerous neural architectures (or models), and return the best matching ones. As for any semantic search engine, an indexed database of all searched architecture embeddings is needed, within which the architecture search is performed. For the search procedure over such database using ArchBERT, the text query is encoded by the text encoder, and then is cross-encoded to make sure the previously-learned architectural knowledge is also utilized for computing final text embeddings. The pooled text

Table 4: Qualitative results on various tasks (✓: Correct/Similar, ✗: Incorrect/Dissimilar, \*: wrong preds).

Architecture	Text	AR	ACD
ResNet18	image classifier with residual layers	✓	✗
Fasterrcnn (ResNet50)	text classifier using bert-based models	✗	
Bert-base	object detection for photos	✗	✓
RoBERT (small)	text classifier using bert-based models	✓	
Vit_b_16	bert-like image classification	✗*	✗
Fasterrcnn (mobilenet)	object detection for photos	✓	
ConvNext (tiny)	a very large convnext architecture	✓*	✗
Bert-mini	language model with attention layers	✓	
Figure 3’s right architecture (AutoNet)	<p><b>AC:</b> "this model separable convolution which divides a single convolution into two convolutions"</p> <hr/> <p><b>AQA:</b> What type of pooling is used in this architecture?</p> <p><b>Prediction:</b> 'MaxPool2d', 'AvgPool2d'</p>		

embeddings are then compared with all the architecture embeddings stored in the database to find the best matching (most similar) architectures. We did not report any numerical analysis for AS due to the lack of related validation set. However, qualitative demo is available in the supplementary materials.

## 5.7 Qualitative Results

In Table 4, ArchBERT’s predictions on AR and ACD tasks over some samples from TVHF validation set are given. In addition, we present the predictions on AC and AQA tasks over the right architecture in Figure 3 (i.e., a sample from AutoNet validation set). Sample cases for which ArchBERT makes wrong predictions are also given in the table (marked with \*), e.g., AR’s prediction for Vit\_b\_16 and ConvNext-tiny architectures.

## 5.8 Ablation Study

We conduct ablation study to analyze the effect of ArchBERT’s different modules such as MAM, Cross Encoder, and graph elements on the performance of AR, ACD, AQA, and AC tasks. The results are summarized in Tables 2 and 3.

First, we remove the MAM head and its loss from the pre-training and fine-tuning stages. The performance of the pre-trained model without MAM is evaluated on AR and ACD with the TVHF dataset. As seen in Table 2, excluding MAM in pre-training results in a significant F1 drops by 7.59%



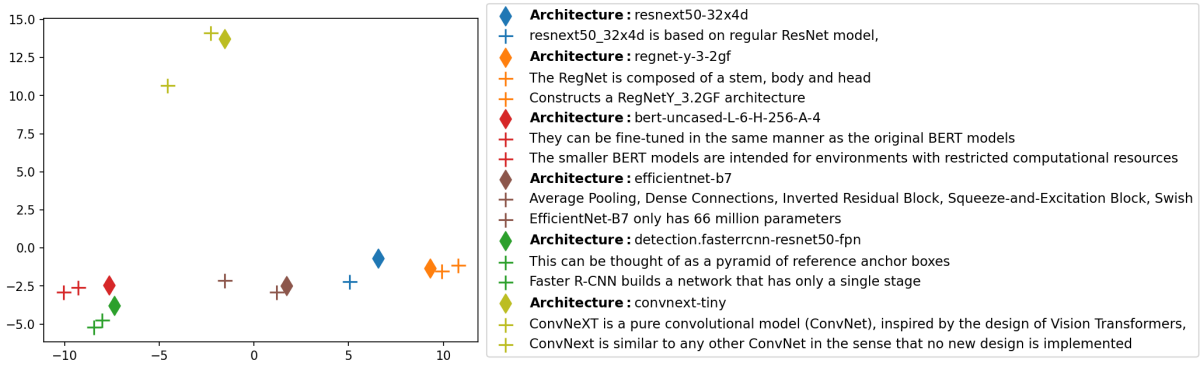


Figure 5: Visualization of example relevant architecture and text embeddings in a 2D space (projected via PCA).

and 10.51% on AR and ACD tasks, respectively. The effect of MAM on finetuned ArchBERT for AQA and AC downstream tasks is also evaluated and reported in in Tables 2 and 3. It is shown that using MAM provides F1 score improvements of 7.35% and 0.03% on AQA and AC, respectively.

We also study the ArchBERT’s performance when the Transformer cross encoder is not used for encoding the architectures. In this case, the embeddings obtained from the architecture encoder are directly used for training and evaluating the model by bypassing the cross encoder. The corresponding results on AR, ACD, and AQA tasks are given in Table 2. From the results, when the cross encoder is removed, the performance of both the pre-trained and fine-tuned models decreases. This reveals the importance of the cross encoder in joint encoding and learning of the text and architecture. As seen in the table, the F1 scores on AR, ACD, and AQA tasks are substantially reduced by 14.83%, 17.75%, and 10.18%, respectively, if the cross encoder is not utilized for architecture encoding.

We also ran a set of ablations over different graph items. For AR, F1 scores of 71.86% (ArchBERT), 69.16% (w/o shape), 68.98% (w/o edge), and 65.80% (w/o shape+edge) are achieved. For ACD, F1 scores of 60.10% (ArchBERT), 60.20% (w/o shape), 47.96% (w/o edge), and 56.45% (w/o shape+edge) are obtained. It is seen that using all graph items provides the best results. For ACD, the shape has no effect on F1 score, but excluding it gives  $\approx 1\%$  lower accuracy.

The ArchBERT’s performance on out-of-distribution data will be presented in the appendix.

### 5.9 Embeddings Visualization

As discussed before, ArchBERT learns to minimize the cosine distance between relevant text and archi-

ture embeddings, while maximizing the distance for the irrelevant ones. To convey this concept, we visualize the joint embeddings of example relevant texts and architectures (i.e.,  $J_t$  and  $J_g$  in Equation 2) from TVHF dataset in Figure 5. The points in the figure are obtained by projecting the embeddings to a 2D space via PCA (Jolliffe, 2005). As shown in Figure 5, the text embeddings are mapped to the points near by their relevant architectures. This implies that ArchBERT has learned to minimize the distance between the related pairs of texts and architectures (i.e., positive samples) and obtain similar embeddings for them. On the other hand, the points for the irrelevant descriptions and architectures are projected far from each other, which shows the success of ArchBERT in maximizing the distance between unrelated pairs.

## 6 Conclusion

In this paper, we proposed ArchBERT, a bi-modal solution for joint learning of neural architectures and natural languages. We also introduced a new pre-training technique called Masked Architecture Modeling (MAM) for a better generalization of ArchBERT. In addition, two new bi-modal benchmark datasets called TVHF and AutoNet were presented on which the proposed model was trained and evaluated for different downstream tasks. Five architecture-language-related tasks and applications were introduced in this work to verify the performance of ArchBERT. This work has opened up new avenues for research in the area of architecture-language joint understanding, particularly the proposed benchmarks. Potential research directions to this work include text-based neural architecture generation and bi-modal learning of languages and other graph-structured modalities such as knowledge graphs and social network graphs.

## References

- Hassan Akbari, Liangzhe Yuan, Rui Qian, Wei-Hong Chuang, Shih-Fu Chang, Yin Cui, and Boqing Gong. 2021. VATT: Transformers for multimodal self-supervised learning from raw video, audio and text. *Advances in Neural Information Processing Systems*, 34:24206–24221.
- Mohammad Akbari, Amin Banitalebi-Dehkordi, and Yong Zhang. 2022. E-lang: Energy-based joint inferencing of super and swift language models. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5229–5244.
- Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6077–6086.
- Tadas Baltrušaitis, Chaitanya Ahuja, and Louis-Philippe Morency. 2018. Multimodal machine learning: A survey and taxonomy. *IEEE transactions on pattern analysis and machine intelligence*, 41(2):423–443.
- Hangbo Bao, Wenhui Wang, Li Dong, and Furu Wei. 2022. VI-beit: Generative vision-language pretraining. *arXiv preprint arXiv:2206.01127*.
- 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.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186.
- Zhangyin Feng, Daya Guo, Duyu Tang, Nan Duan, Xiaocheng Feng, Ming Gong, Linjun Shou, Bing Qin, Ting Liu, Daxin Jiang, and Ming Zhou. 2020. Codebert: A pre-trained model for programming and natural languages. In *Findings of the Association for Computational Linguistics: EMNLP 2020, Online Event, 16-20 November 2020*, volume EMNLP 2020 of *Findings of ACL*, pages 1536–1547. Association for Computational Linguistics.
- Tsu-Jui Fu, Linjie Li, Zhe Gan, Kevin Lin, William Yang Wang, Lijuan Wang, and Zicheng Liu. 2021. Violet: End-to-end video-language transformers with masked visual-token modeling. *arXiv preprint arXiv:2111.12681*.
- Google. 2022. The size and quality of a data set. <https://developers.google.com/machine-learning/data-prep/construct/collect/data-size-quality>, Last accessed on 2022-12-14.
- Daya Guo, Shuo Ren, Shuai Lu, Zhangyin Feng, Duyu Tang, Shujie Liu, Long Zhou, Nan Duan, Alexey Svyatkovskiy, Shengyu Fu, Michele Tufano, Shao Kun Deng, Colin B. Clement, Dawn Drain, Neel Sundaresan, Jian Yin, Daxin Jiang, and Ming Zhou. 2021. Graphcodebert: Pre-training code representations with data flow. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Matthew Henderson, Rami Al-Rfou, Brian Strope, Yun-Hsuan Sung, László Lukács, Ruiqi Guo, Sanjiv Kumar, Balint Miklos, and Ray Kurzweil. 2017. Efficient natural language response suggestion for smart reply. *arXiv preprint arXiv:1705.00652*.
- Danfeng Hong, Lianru Gao, Naoto Yokoya, Jing Yao, Jocelyn Chanussot, Qian Du, and Bing Zhang. 2020. More diverse means better: Multimodal deep learning meets remote-sensing imagery classification. *IEEE Transactions on Geoscience and Remote Sensing*, 59(5):4340–4354.
- Ian Jolliffe. 2005. Principal component analysis. *Encyclopedia of statistics in behavioral science*.
- Boris Knyazev, Michal Drozdal, Graham W Taylor, and Adriana Romero Soriano. 2021. Parameter prediction for unseen deep architectures. *Advances in Neural Information Processing Systems*, 34:29433–29448.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2017. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880.
- Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2019. VisualBERT: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. ViLBERT: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *Advances in neural information processing systems*, 32.
- Maitiniyazi Maimaitijiang, Vasit Sagan, Paheding Sidike, Sean Hartling, Flavio Esposito, and Felix B Fritsch. 2020. Soybean yield prediction from uav using multimodal data fusion and deep learning. *Remote sensing of environment*, 237:111599.

- Sébastien Marcel and Yann Rodriguez. 2010. Torchvision the machine-vision package of torch. In *Proceedings of the 18th ACM International Conference on Multimedia*, page 1485–1488. Association for Computing Machinery.
- Daniel Michelsanti, Zheng-Hua Tan, Shi-Xiong Zhang, Yong Xu, Meng Yu, Dong Yu, and Jesper Jensen. 2021. An overview of deep-learning-based audio-visual speech enhancement and separation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:1368–1396.
- OpenAI. 2022. Introducing chatgpt. <https://openai.com/blog/chatgpt>.
- OpenAI. 2023. [Gpt-4 technical report](#).
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67.
- Nils Reimers and Iryna Gurevych. 2019. Sentencebert: Sentence embeddings using siamese bert networks. 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 3982–3992.
- Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28.
- Julio Santisteban and Javier Tejada-Cárcamo. 2015. Unilateral jaccard similarity coefficient. In *GSB@SIGIR*, pages 23–27.
- Liam Schoneveld, Alice Othmani, and Hazem Abdelkawy. 2021. Leveraging recent advances in deep learning for audio-visual emotion recognition. *Pattern Recognition Letters*, 146:1–7.
- Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. 2019. VideoBERT: A joint model for video and language representation learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7464–7473.
- Hao Tan and Mohit Bansal. 2019. LXMERT: learning cross-modality encoder representations from transformers. 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 2019, Hong Kong, China, November 3-7, 2019*, pages 5099–5110. Association for Computational Linguistics.
- Yao-Hung Hubert Tsai, Shaojie Bai, Paul Pu Liang, J Zico Kolter, Louis-Philippe Morency, and Ruslan Salakhutdinov. 2019. Multimodal transformer for unaligned multimodal language sequences. In *Proceedings of the conference. Association for Computational Linguistics. Meeting*, volume 2019, page 6558. NIH Public Access.
- Luís A Vale-Silva and Karl Rohr. 2021. Long-term cancer survival prediction using multimodal deep learning. *Scientific Reports*, 11(1):1–12.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph attention networks. In *International Conference on Learning Representations*.
- Janani Venugopalan, Li Tong, Hamid Reza Hassanzadeh, and May D Wang. 2021. Multimodal deep learning models for early detection of alzheimer’s disease stage. *Scientific reports*, 11(1):1–13.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.
- Yi Xiao, Felipe Codevilla, Akhil Gurrum, Onay Urfalioglu, and Antonio M López. 2020. Multimodal end-to-end autonomous driving. *IEEE Transactions on Intelligent Transportation Systems*.

## A Appendix

### A.1 Code, Dataset, and Demo

In order for the results to be reproducible, we share our test code (plus the pre-trained model files) with detailed instructions in the supplementary materials. The code also includes the scripts for generating both TVHF and AutoNet datasets.

We also uploaded 6 video files demonstrating the performance of ArchBERT on the following downstream tasks: architecture search (AS), architectural reasoning (AR), architecture clone detection (ACD), bi-modal architecture clone detection (BACD), architectural question answering (AQA), and architecture captioning (AC).

All the code and demo files are also available [here](#)<sup>2</sup>.

BACD task is similar to ACD, except that a supporting text, which is considered as an extra criteria to refine the results, is also provided along with the two given architectures. The average similarity of the architectures' embeddings with the help of the text embeddings is evaluated to check if the architectures are similar or not.

The video recordings were taken from a web application we built to demonstrate the real-world application of our method. Example screenshots of the AR and BACD demos are shown in Figure 6.

### A.2 ArchBERT's Performance on OOD Data

In order to study the behaviour of ArchBERT on out-of-distribution (OOD) data, we establish another set of experiments on individual TV and HF datasets that have different distributions. In this regard, we pre-train ArchBERT on each of TVHF,

<sup>2</sup><https://developer.huaweicloud.com/develop/aigallery/notebook/detail?id=e6a924c7-735a-4e02-a25b-4416b77b6315>

The screenshot shows a web interface for Architectural Reasoning (AR). It features a text input field containing "image classifier with residual layer" with a cursor at the end. Below the input is a dropdown menu for "Input architecture" showing "aa8471ad-62e8-4ad2-b234-ffe9add34bb9\_\_\_resnet18". The "Dataset" is set to "TVHF" and "#Samples" is "100000". At the bottom, it states "The statement is: Correct."

TV-only, and HF-only datasets, and evaluate their performance on each other. The corresponding experimental results are summarized in Table 5.

As observed in the table, the models trained on TV and HF subsets do not generalize to each other due to the difference in their data distributions, which results in poor performance. The distribution plots for TV and HF subsets are shown in Figure 8. As given in Table 5, the highest scores on each of TV and HF subsets are obtained by the model trained with the entire TVHF training dataset. In order to improve the performance of our model on OOD, some techniques such as zero-shot or few-shot learning can be employed, which is a potential research direction for this work.

### A.3 Embeddings Visualization

In Figure 5, an embedding visualization of some architecture-text pairs was illustrated. In Figure 7, the visualizations for two different architectures from TVHF dataset are individually presented. The points on the figures are obtained by projecting the final ArchBERT's embeddings onto a 2D space via PCA. As shown in the plots, unlike the relevant text embeddings (marked with +), the irrelevant ones (marked with ×) are projected far from the corresponding architecture embeddings.

### A.4 Data Generation

The procedure of creating TVHF dataset along with negative samples are given in Algorithm 1. To generate the negative data samples, a pre-trained S-BERT model (Reimers and Gurevych, 2019) is used to calculate the similarity score between all possible pairs of unique descriptions. If the maximum similarity score between each unique sentence and all other sentences of each unique neural architecture is smaller than a threshold 0.5, that sen-

The screenshot shows a web interface for Bi-Modal Architecture Clone Detection (BACD). It features a "Query" input field containing "model for object detector" with a cursor. Below it are two dropdown menus for "Input architecture 1" (showing "86bac9ef-9591-4874-be1a-919f65410724\_\_\_detection.ssd300\_vgg16") and "Input architecture 2" (showing "676af582-2358-4d40-9b62-1c7088a72910\_\_\_detection.fasterrcnn\_resnet50\_fpn"). The "Dataset" is "TVHF" and "#Samples" is "100000". At the bottom, it states "The architectures are: Similar."

Figure 6: Screenshots from the demos. Left: Architectural Reasoning (AR); Right: Bi-Modal Architecture Clone Detection (BACD).

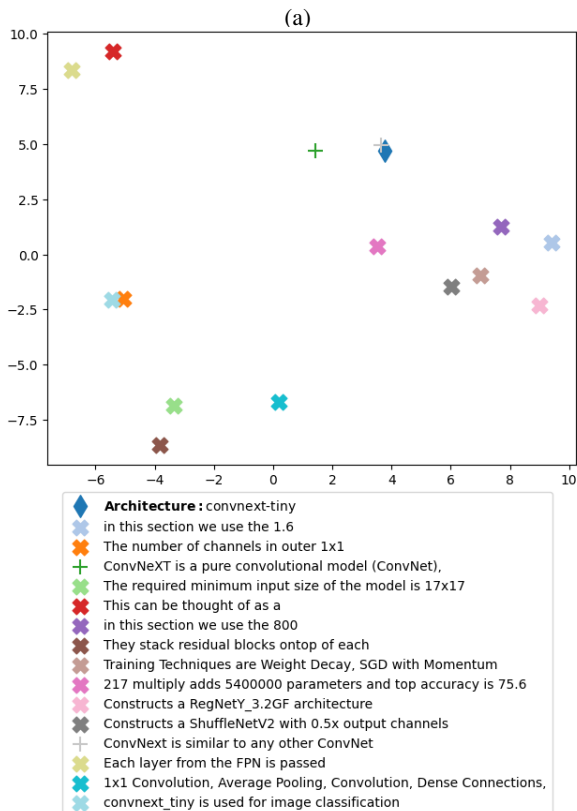
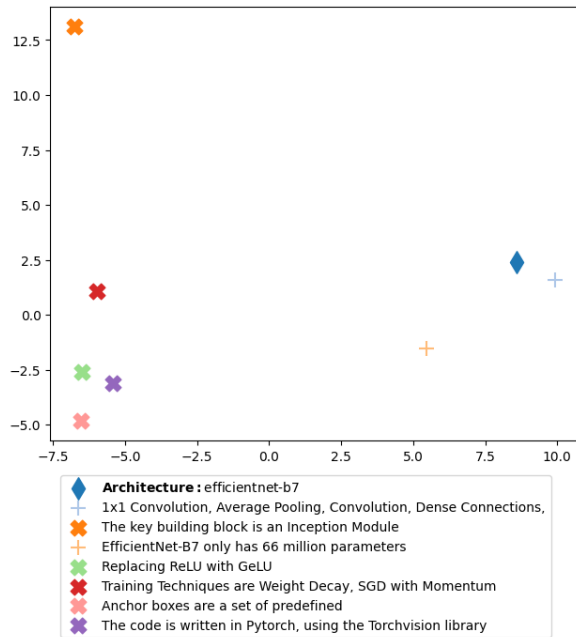


Figure 7: Visualization of example pairs of (ir)relevant architecture and text embeddings in a 2D space (projected via PCA).

Train set	Task	F1 on Validation set		
		TV	HF	TVHF
TV	AR	85.05	3.82	28.78
	ACD	58.88	22.85	23.30
HF	AR	9.19	64.26	42.43
	ACD	15.42	59.98	54.57
TVHF	AR	<b>85.32</b>	<b>64.39</b>	<b>71.86</b>
	ACD	<b>62.77</b>	<b>60.01</b>	<b>60.10</b>

Table 5: ArchBERT’s performance on OOD data.

tence is chosen as an irrelevant description for that specific neural architecture. Note that 93% of the final TVHF train set contains negative samples. The above-mentioned procedure of generating many negative candidates per each positive sample was inspired by the multiple negatives sampling idea described by Henderson et al. (2017). Having multiple negatives was proved to be effective when used with dot-product and cosine similarity loss function (Equation 3 in the main paper).

For TVHF-ACD dataset, all possible pairs of neural architectures were compared based on their structures. A hard score of 1 or 0 is then assigned to a similar or dissimilar pair of architectures, respectively. For TorchVision architectures with the same architectural base (e.g., ResNet family), a hard score of 1 is assigned to the pair. For HuggingFace models, the configuration files were compared and in case of having similar specifications, a hard score of 1 has been assigned to those architectures. In overall, the TVHF-ACD dataset includes 11% of similar pairs of architectures.

For AutoNet dataset, all unique layers of each architecture are first extracted. To do so, an algorithm is developed to take an architecture as input and recursively extracts all unique modules and their class path within that architecture. These unique layers are then used along with a list of various pre-defined templates to randomly generate meaningful descriptions with different words and sentence structures. The algorithm is then used with modules that are not included in the architecture to generate irrelevant descriptions that are considered as negative data samples. Each architecture has about 10-11 different descriptions about 30% of which are the positive ones. The same extracted layers and procedures are also used for automati-

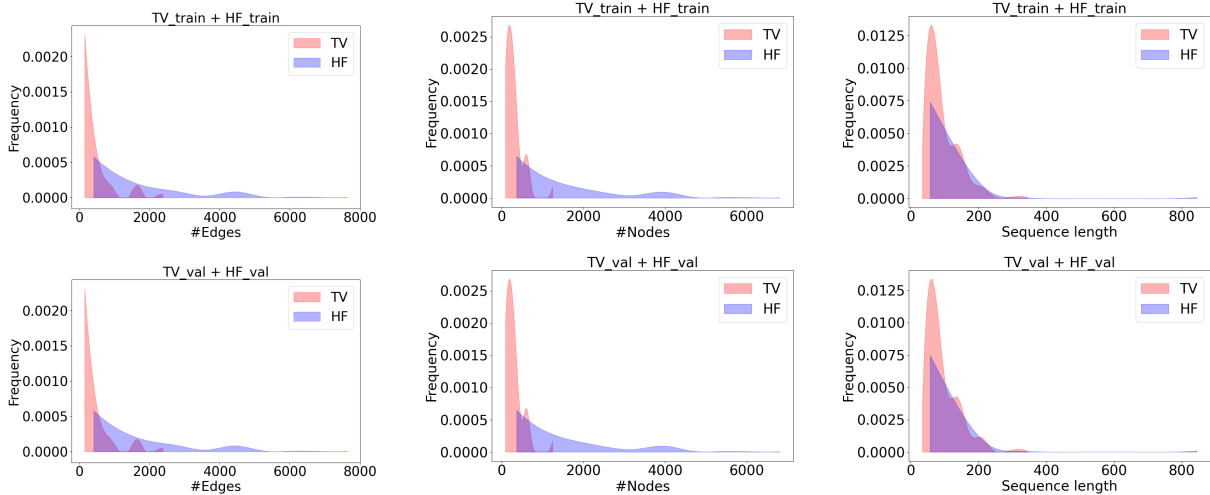


Figure 8: Distribution plots of TV and HF train and validation sub-datasets compared with each other.

cally generating the question and answer pairs, but with a different set of templates for questions.

---

#### Algorithm 1 TVHF dataset generator

---

**Input:** Threshold  $\beta$ , architectures  $G$ , pos\_samples  $T^p$   
**Output:** list of architectures plus their positive and negative descriptions  
**for** each unique neural architecture  $G_j \in G$  **do**  
  **for** each unique description  $T_i^p \in T^p(G_j)$  **do**  
    **if**  $\max(\text{SBERT}(T_i^p, T_{\sim i}^p)) \leq \beta$  **then**  
      Add  $T_i$  to  $T^n(G_j)$  (list of neg\_samples for  $j$ th architecture)  
    **end if**  
  **end for**  
**end for**  
**return**  $\{G, (T^p, T^n)\}$

---

### A.5 Distribution Plots for TVHF and AutoNet

Figure 9 shows the distribution plots of the TVHF, AutoNet, and AutoNet-AQA datasets. For each dataset, the plots of the training and validation distributions of the number of nodes, the number of edges, the number of textual tokens, and the sequence length of the descriptions are illustrated.

### A.6 Sample Data from TVHF and AutoNet

In Table 6, example positive architecture-description pairs (for both computer vision and natural language processing problems) from TVHF dataset are given.

Some sample pairs of architectures (with their corresponding "similar" or "dissimilar" ground

truth labels) from TVHF-ACD dataset are also presented in Table 7.

In Table 9, we also provide data samples for the BACD task, which includes quartets of two architectures, supporting description, and the similarity label. Note that the numerical analysis of ArchBERT over BACD is not provided because our BACD validation dataset is not finalized to be used for this matter.

Table 8 also presents a few data samples from AutoNet dataset used for fine-tuning and evaluating ArchBERT on AC task. In Table 10, sample data from AutoNet-AQA including the automatically generated questions and ground truth answers for AQA downstream task are given.

In Figures 10 and 11, the visualization of all graphs generated for the neural architectures listed in Tables 4, 8, and 10 are illustrated.

### A.7 Dataset Quality Analysis

We provide dataset quality analysis based on four criteria: reliability and completeness, label/feature noise, feature representation, and minimizing skew (Google, 2022).

#### A.7.1 Reliability and Completeness

The reliability of data refers to how trustable the data is, whether it has duplicated values and if it covers both positive and negative samples. As for dataset completeness, it refers to how much of the relevant information is included in the dataset for dealing with the desired problem.

In our TVHF dataset, we have collected models and their relevant descriptions as related bi-modal data types for the ArchBERT model to learn neu-

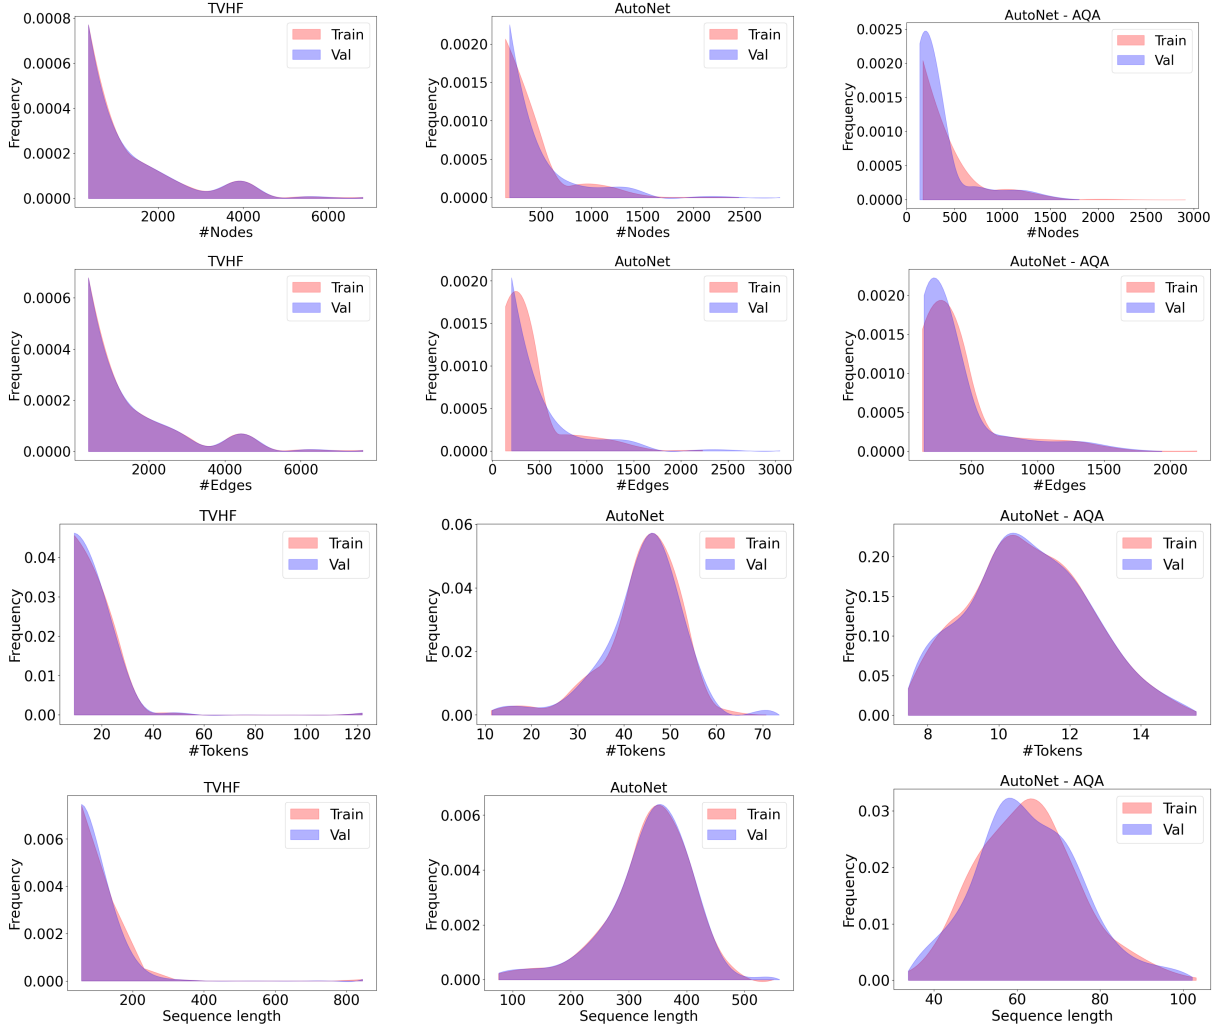


Figure 9: Distribution plots of TVHF, AutoNet, and AutoNet-AQA train/validation datasets.

ral architectures along with their corresponding natural language descriptions. We considered the reliability and completeness of our dataset by collecting various models with different architectures designed for different tasks such as image and text classification, object detection, text summarization, etc. Also, the descriptions that have been assigned to each model were collected through blog posts, articles, papers, and documentations containing both high/low-level information related to that specific model. Due to the limited number of human-designed models, to make our dataset large enough for training purposes, we used each architecture more than once, and each time we assigned a different unique description to it to avoid having duplicate architecture-description pairs in our dataset. Moreover, we generated negative samples by assigning irrelevant descriptions to the architectures, so that the model could learn both similarities and dissimilarities.

As discussed in Section 4, some of the descriptions in TVHF dataset did not include relevant technical information to the corresponding models. We manually reviewed the descriptions and removed such samples. We will further enhance the descriptions associated with each model within the release of the next version of our dataset.

### A.7.2 Label/Feature Noise

Label noise refers to an imperfect annotation of data that confounds the assessment of model performance when training machine learning models. Feature noise can be defined as the noise got into the dataset through various factors such as incorrect collection by humans or instruments. Inconsistencies in data formats, missing values, and outliers are examples of noise created by this process.

If noise in a dataset is defined as a wrong description for a model, our dataset is a noise-free dataset because we annotated the samples manually.

Since the description of building blocks in the AutoNet models are converted to textual descriptions and question samples automatically, all the generated samples are relevant and noise-free.

For our ACD dataset, we manually hard-labeled the models based on their similarity with each another. Therefore, there is no missed or wrongly labeled example in the entire dataset.

### **A.7.3 Feature Representation**

Mapping data to useful features while presenting them to the model is defined as feature representation. In this case, we consider how data is presented to the model and whether the numeric values need to be normalized.

To show our data to the ArchBERT model, we have been consistent in the following way. For architectures, based on their computational graphs, we extracted nodes, shapes, and edges, which the major and sufficient items to represent an architecture in our work. We then normalized these items and passed them to the model. As for descriptions, we represented each textual description with tokens, normalized them, and used them as inputs to the model.

### **A.7.4 Minimizing Skew**

One of the reasons that may cause getting different results for computed metrics at training vs. validation stages is training/validation skew. It usually happens when different features are presented to the model in training and validation stages.

We have collected our data and presented them to the model in the way that both training and validation stages receive the exact same set of features coming from the same distribution. This guarantees that our data is not skewed towards training or validation stages.



Table 6: Positive data samples from the TVHF dataset (TV: TorchVision, HF: HuggingFace).

Architecture	Description	Source
vit_b_16	adopted from BERT	TV
segmentation.deeplabv3_resnet101	Improved version of DeepLab v2, with optimization of ASPP layer hyper parameters and without a Dense CRF layer, for faster operation.	TV
resnet101	Residual Networks, or ResNets, learn residual functions with reference to the layer inputs , instead of learning unreferenced functions.	TV
densenet121	A DenseNet is a type of convolutional neural network that utilises dense connections between layers, through Dense Blocks, where we connect all layers (with matching feature-map sizes) directly with each other	TV
resnext50_32x4d	ResNeXt is a homogeneous neural network which reduces the number of hyper parameters required by conventional ResNet.	TV
detection.keypointcnn_resnet50_fpn	12 Million Parameters, 2 Billion FLOPs and File Size is 47.08 MB.	TV
DemangeJeremy/4-sentiments-with-flaubert	This model is a fine-tuned version of google/fnet-base on the GLUE WNLI dataset	HF
ctoraman/RobERTa-TR-medium-char	Model architecture is similar to bert-medium (8 layers, 8 heads, and 512 hidden size)	HF
google/t5-efficient-base-dm1000	T5-Efficient-BASE-DM1000 is a variation of Google’s original T5 following the T5 model architecture.	HF
microsoft/unihanlm-base	a self-supervised Chinese-Japanese pre-trained masked language model (MLM) with a novel two-stage coarse-to-fine training approach.	HF
facebook/wmt21-dense-24-wide-en-x	WMT 21 En-X is a 4.7B multilingual encoder-decoder (seq-to-seq) model trained for one-to-many multilingual translation.	HF

Table 7: Positive and negative data samples from TVHF-ACD validation set (TV: TorchVision, HF: HuggingFace, 0: dissimilar, 1: similar).

Architecture 1	Architecture 2	Label	Source
vgg11	vgg19_bn	1	TV
mnasnet0_5	mnasnet0_75	1	TV
inception_v3	efficientnet_b3	0	TV
efficientnet_b1	regnet_x_800mf	0	TV
google/t5-efficient-large-kv128	google/t5-efficient-small-kv16	1	HF
jweb/japanese-boseki-gpt2-1b	tartuNLP/gpt-4-est-large	1	HF
hakurei/gpt-j-random-tinier	minimaxir/magic-the-gathering	0	HF
mwesner/bart-mlm	tartuNLP/gpt-4-est-base	0	HF

Table 8: Positive and negative data samples from AutoNet (Architecture: list of unique layers, 0: negative sample, 1: positive sample).

Architecture	Description	Label
'Conv2d', 'PosEnc', 'ReLU',	This architecture contains 2d max pooling layer which is a pooling operation that calculates the maximum value, and Gaussian Error Linear Units (gelu) activation function which is a smoother version of RELU. It also has 2D Adaptive Average pooling layer.	1
'BatchNorm2d', 'Linear', 'Dropout', 'LayerNorm', 'GELU',	This neural network has Layer normalization over input across the features instead of batch dimension, and linear module which applies a linear transformation to the incoming data. It also contains Dropout layer that is used to drastically reduce the chance of overfitting during training.	1
'Dil_conv2d', 'Zero', 'MaxPool2d', 'AvgPool2d', 'AdaptiveAvgPool2d'	This classification neural network includes 2D average pooling layer used for calculating the average for each patch of the feature map and has about 1.18 Million parameters. In Total, this neural network architecture has 432 layers, and, it has 95 Conv2d, 1 PosEnc, 80 ReLU, 79 BatchNorm2d, 62 Linear, 46 Dropout, 30 LayerNorm, 15 GELU, 15 Dil_conv2d, 4 Zero, 2 MaxPool2d, 2 AvgPool2d, and 1 AdaptiveAvgPool2d layer.	1
	This neural architecture has 2D frozen batch normalization module in which the batch statistics and the affine parameters are fixed, and Anchor Generator module which is a standard for 2D anchor-based detectors. Additionally, this architecture contains stochastic depth layer which aims to shrink the depth of a network during training.	0
	This classifier includes 2D transposed convolution layer that applies convolution with a fractional stride.	0
'Conv2d', 'Hardswish', 'GeLU', 'AvgPool2d', 'Sep_conv2d', 'AdaptiveAvgPool2d', 'Dropout'	This classification neural architecture has Separable Convolution which divides a single convolution into two or more convolutions to reduce the number of parameters while producing the same output, and Hard Swish activation function that replaces the computationally expensive sigmoid with a piecewise linear analogue. This classifier also includes 2D average pooling layer used for calculating the average for each patch of the feature map.	1
	This network includes Dropout layer that is used to drastically reduce the chance of overfitting during training, and 2D Adaptive Average pooling layer. This neural architecture has about 0.38 Million parameters.	1
	This classification architecture includes generalized rnn transform module which performs input transformation before feeding the data to a GeneralizedRCNN model, and Quantize stub module that is a place holder for quantize operation. Another part of this neural network is ReLU6 activation function which is a modification of the rectified linear unit (relu) where the activation is limited to a maximum size of 6.	0
	This architecture contains Layer normalization over input across the features instead of batch dimension, and dequantization module which converts tensors from quantized to floating point.	0

Table 9: Positive and negative data samples for BACD task (TV: TorchVision, HF: HuggingFace, 0: dissimilar, 1: similar).

Architecture 1	Architecture 2	Supporting text	Label	Source
resnet18	segmentation.fcn_resnet101	A model containing residual connection	1	TV
mnasnet0_5	vgg19	An architecture for image classification	1	TV
wide_resnet101_2	segmentation.deeplabv3_resnet50	An architecture for image classification	0	TV
resnet34	alexnet	A model containing residual connection	0	TV
ctoraman/ RoBERTa-TR-medium-char	ctoraman/ RoBERTa-TR-medium-wp-66k	Model architecture is similar to bert-medium	1	HF
dbmdz/ electra-base-turkish-cased-discriminator	skplanet/ dialog-koelectra-small-generator	containing ELECTRA for self-supervised language representation learning	1	HF
rmihaylov/ pegasus-base-cnn-dailymail-bg	TristanBehrens/js-fakes-4bars	A model for summarization	0	HF
facebook/ m2m100-12B-avg-10-ckpt	google/t5-11b-ssm-nqo	A pre-trained model for Question Answering	0	HF

Table 10: Data samples from AutoNet-AQA (Architecture: list of unique layers).

Architecture	Question	Ground Truth Answer
Conv2d, BatchNorm2d, ReLU, Dil_conv2d, Sep_conv2d, AvgPool2d, AdaptiveAvgPool2d, Linear	what type of pooling module has been used in this neural architecture?	AvgPool2d, AdaptiveAvgPool2d
	what 2d average pooling layer performs in this neural network?	calculating the average for each patch of the feature map
	what 2d Dilated Convolution module does in this network?	creating a wider kernel by inserting spaces between the kernel elements
	what 2d max pool kernel size has been used in this network?	This model does not include MaxPool2d
	in general what kernel size are used in this neural network model?	5*5,1*1,3*3
'Conv2d', 'GELU', 'MaxPool2d', 'LayerNorm', 'Linear', 'Hardswish', 'Dil_conv2d', 'LayerNorm'	what 2d max pooling module calculates in this neural network?	calculating the maximum value for each patch of the feature map
	what type of normalization layer is used after convolution in this neural network architecture?	LayerNorm
	what type of activation layer has been used in this neural network model?	GELU, Hardswish
	what hard sigmoid module performs in this model?	This model does not include Hardsigmoid
	overall what kind of layers are included in this neural network architecture?	'Conv2d', 'GELU', 'MaxPool2d', 'LayerNorm', 'Linear', 'Hardswish', 'Dil_conv2d', 'LayerNorm'

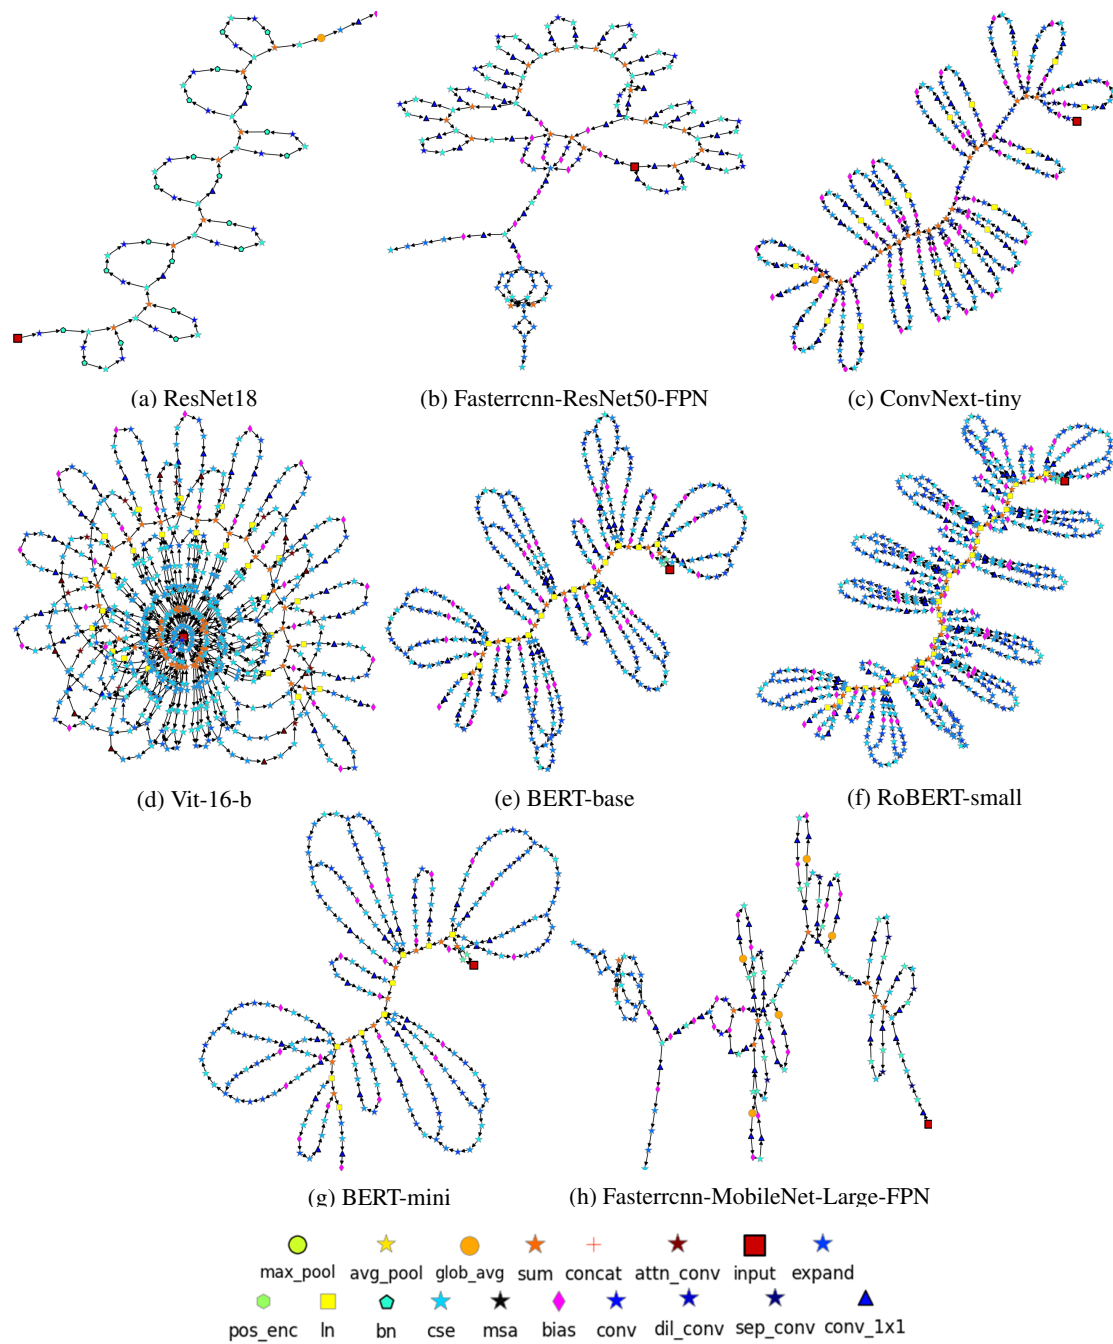
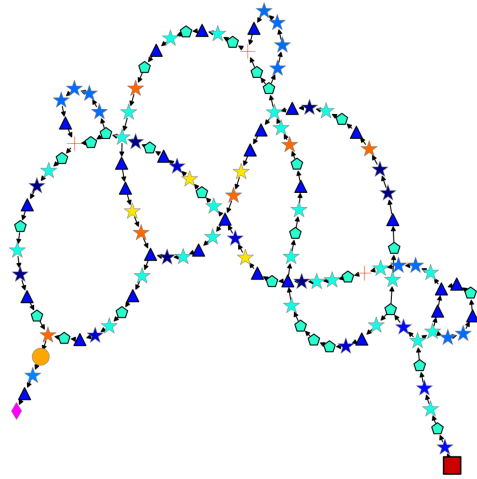
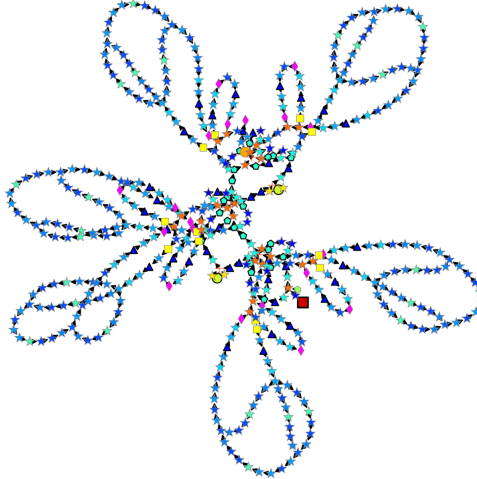


Figure 10: Graphs generated for the architectures listed in Table 4



(a) Architecture with layers: Conv2d, Batch-Norm2d, ReLU, Dil\_conv2d, Sep\_conv2d, Avg-Pool2d, AdaptiveAvgPool2d, Linear



(b) Architecture with layers: Conv2d, PosEnc, ReLU, BatchNorm2d, Linear, Dropout, LayerNorm, GELU, Dil\_conv2d, Zero, MaxPool2d, AvgPool2d, AdaptiveAvgPool2d



Figure 11: Graphs generated for the architectures listed in Tables 8 and 10.