

Automatically Generating Dockerfiles via Deep-Learning: Challenges and Promises

Giovanni Rosa*, Antonio Mastropaolo†, Simone Scalabrino*, Gabriele Bavota†, Rocco Oliveto*

*STAKE Lab - University of Molise, Pesche, Italy

{name.surname}@unimol.it

†Software Institute - Università della Svizzera Italiana (USI), Switzerland

{name.surname}@usi.ch

Abstract—Containerization allows developers to define the execution environment in which their software needs to be installed. Docker is the leading platform in this field, and developers that use it are required to write a Dockerfile for their software. Writing Dockerfiles is far from trivial, especially when the system has unusual requirements as for its execution environment. Despite several tools exist to support developers in writing Dockerfiles, none of them is able to generate entire Dockerfiles from scratch given a high-level specification of the requirements of the execution environment. In this paper, we present a study in which we aim at understanding to what extent Deep Learning (DL), which has been proven successful for other coding tasks, can be used for this specific coding task. We preliminarily defined a structured natural language specification for Dockerfile requirements, and a methodology that we use to automatically infer the requirements from the largest dataset of Dockerfiles currently available. We used the obtained dataset, with 670,982 instances, to train and test a Text-to-Text Transfer Transformer (T5) model, following the current state-of-the-art procedure for coding tasks, to automatically generate Dockerfiles from the structured specifications. The results of our evaluation show that T5 performs similarly to more trivial IR-based baselines we considered. We also provide some guidance for you to explore to try and successfully harness the power of deep learning models for Dockerfile generation.

Index Terms—docker, deep learning

I. INTRODUCTION

Software companies are more and more often starting adopting the DevOps methodology for developing their products. DevOps strongly relies on technologies for automating the build and the deployment of the systems (CI/CD), and it results in shorter release cycles [1]. In this context, containerization technologies are fundamental to allow developers take control over the execution environment of their products. Such tools allow developers to reduce the risks of issues arising from possible differences between the development/testing environment and the production environment. Docker is the leading containerization technology, becoming the “Most Loved” and “Most Wanted” platform, according to the 2022 Stack Overflow survey [2]. When using Docker for a specific software product, developers are required to write a Dockerfile, which contains a sequence of instructions that, when executed, allow to build an image (Docker image) which can be run in one or more lightweight virtual machines (Docker containers).

Writing Dockerfiles is not trivial. First, basic system administration skills are required, and developers do not always

Fig. 1: Example of Dockerfile for Tomcat and FFMpeg.

```
1 FROM tomcat:7.0.75-jre8
2
3 RUN echo deb http://archive.ubuntu.com/ubuntu precise universe multiverse
   >> /etc/apt/sources.list; apt-get update && \
4 apt-get -y --fix-missing install autoconf automake build-essential \
5 git mercurial cmake libass-dev libgpg-dev libtheora-dev libtool \
6 libvdpau-dev libvorbis-dev pkg-config texi2html zlib1g-dev \
7 libmp3lame-dev wget yasm && \
8 apt-get clean
9
10 WORKDIR /usr/local/src
11 # install x265
12 RUN hg clone https://bitbucket.org/multicoreware/x265 && \
13 cd /usr/local/src/x265/build/linux && \
14 cmake -DCMAKE_INSTALL_PREFIX=PATH=/usr ../../source && \
15 make -j 8 && \
16 make install
17
18 WORKDIR /usr/local/src
19 # install ffmpeg
20 RUN git clone --depth 1 && \
21 cd ffmpeg && \
22 git://source.ffmpeg.org/ffmpeg && ./configure \
23 --extra-libs="ldl" --enable-gpl --enable-libass \
24 --enable-libvorbis --enable-libx265 --enable-nonfree && \
25 make -j 8 && \
26 make install
27
28 WORKDIR /
```

have them. Second, while basic Dockerfiles templates can be used as a starting point, they need to be adapted to the specific requirements of the software system at hand. To clarify this, let us consider the scenario in which a developer needs an environment that contains both Apache Tomcat as a web server and FFMpeg for processing videos, with the support of the x265 codec. While templates exist for the former, they do not contain hints about how to provide also the latter. Fig. 1 shows an example of possible solution: Such a Dockerfile starts from a pre-defined image which contains Tomcat and it installs FFMpeg with the support of x265 on top of it. It can be noticed that determining the commands required to achieve the latter goal requires a moderate effort and it is prone to errors. This problem has been observed in the literature. For example, in the survey conducted by Reis *et al.* [3], the authors observed that developers (especially the less experienced ones) perceive the creation of Dockerfiles as a time consuming activity. Moreover, using Dockerfiles from tutorials and blob posts as a support for their creation can lead to broken Docker images [4].

Some tools and approaches have been introduced to help developers to write Dockerfiles. Some of them [5]–[9] take as input the context (*i.e.*, the root folder of the project) and suggest whole Dockerfiles based on it. Such tools are handy since they allow developers to quickly have a base from which they can start. They also require limited effort from

the developers (*i.e.*, they only need to specify the root of the project). However, they are not able to recognize the need for specific libraries (like FFMpeg, in the previous example). Other tools provide a more in-depth support, but they are limited to specific programming languages (*e.g.*, DockerizeMe for Python). The approach recently introduced by Ye *et al.* [10] can recommend packages that need to be installed given an initial set of dependencies, but it only supports developers in the definition of the dependencies to install. The developer still needs to write complex instructions, like the ones required to build the FFMpeg and x265 libraries in Fig. 1.

Finally, there are code completion tools that support developers while writing Dockerfiles. A recent example is GitHub Copilot [11], which relies on Deep Learning (DL) to achieve this goal. Such tools, however, require that developers manually write part of the Dockerfile, so that they can complete it. None of such tools and approaches is able to generate complete Dockerfiles from a high-level description of what the developer wants in the Dockerfile (*i.e.*, requirements). Previous research [12]–[15] show that DL is a viable solution for code generation-related tasks. However, to the best of our knowledge, no previous work tested to what extent DL can be used to generate complete Dockerfiles.

In this paper, we aim to fill this gap. We first define the format of a structured high-level specification to define, via natural language, the requirements for the definition of Dockerfiles. Then, we define a methodology for automatically inferring such a high-level specification (HLS) from existing Dockerfiles, so that we can build a dataset large enough to train and test a DL model. To this aim, we rely on the largest collection of Dockerfiles available in the literature [16], containing 9.4M Dockerfile snapshots extracted from **all** the open-source projects hosted on GitHub. We run our specification-inference tool on them and, after a filtering procedure, we end up with a set of 670,982 pairs (HLS, Dockerfile). We use this dataset to train and test a state-of-the-art DL model, the Text-to-Text Transfer Transformer (T5) [17], which has been proven effective when supporting several coding tasks [13], [14], following the same pipeline defined in the literature. We compare the DL-based approach with two Information Retrieval (IR)-based approaches, and we check to what extent, given a HLS, the output Dockerfiles of the three techniques: (i) meet the input requirements, (ii) are similar to the target Dockerfile, and (iii) allow to build a Docker image similar to the target one.

We obtain mixed results: While T5 achieves similar results to the best baseline in terms of adherence to the requirements, it generates Dockerfiles less similar to the target Dockerfile. On the other hand, we found that the build of the Dockerfiles generated with T5 succeeds more often and that a higher percentage of intermediate layers produced during the build match the ones obtained with the target Dockerfiles. Interestingly, we also found that T5 truncates the Dockerfiles. This might be due to two main issues. First, **a larger training dataset might be needed for this task**: Despite we consider the largest collection of Dockerfiles in the literature [16],

our results suggest that the T5 learning could benefit from more training data as the existing studies working on regular source code. Second, **a different T5 training stop criterion needs to be defined**: The stop criterion we adopt, which is the one currently used for coding tasks [14], is based on the convergence in terms of BLEU-4 score of the generated predictions. However, considering our results, it seems to be ineffective in the evaluated context.

II. BACKGROUND & RELATED WORK

Writing a Dockerfile is the first step to containerize an application using Docker. A Dockerfile specifies the dependencies and execution environment to build and execute the target application using domain-specific instructions [18]. The first instruction is the FROM, which defines on which existing image the new image builds upon, *i.e.*, the *base image*. Then, we have instructions that perform actions during the build, such as RUN, to execute scripts, WORKDIR, to change the working directory, and COPY/ADD, to copy files, folders and compressed archives. Other instructions are for metadata and documentation (LABEL, EXPOSE), build arguments (ARG), environment variables (ENV), and commands that the container must execute when the image is run (CMD, ENTRYPOINT). Each Docker image built from a Dockerfile is uniquely identified by a *digest*, a hash value computed upon the files that compose the image. While building a Dockerfile, Docker creates a *layer* for each instruction, which constitutes a temporary image that allows to cache and speed-up the next build by avoiding the build of previously-built layers, when possible. The layers are also identified by a unique *digest*.

Several tools and approaches are available to help developers in writing Dockerfiles. Due to space constraints, we only focus on such approaches, which are the most related to our work. *GitHub Copilot* [11] has been recently introduced as a general-purpose code completion tool, and it works also with Dockerfiles. *Humpback* [19] achieves a similar goal, but it is specifically designed for Dockerfiles. The tool by Zhang *et al.* [20] helps developers choosing the right *base image*. Other approaches provide more advanced support. Horton *et al.* [5] proposed DOCKERIZEME, an approach for the automatic inference of environment dependencies starting from Python source code, without requiring inputs from developers. However, it only targets the execution dependencies for Python code. Ye *et al.* [10], proposed DOCKERGEN, an approach that uses knowledge graphs, built upon 220k Dockerfiles, for generating Dockerfiles for a specific software application. Starting from a target software, DOCKERGEN infers all the dependencies required for the execution environment including the selection of a suitable base image. However, the support provided by DOCKERGEN is limited to the recommendation of dependencies, while we aim at generating complete Dockerfiles (including, for example, the RUN instructions required for building an external package, as the ones in Fig. 1). Other tools take as input the context (*i.e.*, the root folder of the project) and suggest whole Dockerfiles based on it. An example is the tool *starter* [7], that generates a Dockerfile and a *docker-*

Field	Type	Description	% Positive Answers
<i>OS</i>	String	Operating system to be used.	75%
<i>Package Manager</i>	String	Linux package manager to be used (<i>e.g.</i> , apt or apk).	58%
<i>Dependencies</i>	String[]	Dependencies (<i>e.g.</i> , packages) that must be provided.	83%
<i>Download of External Dependencies</i>	Boolean	Dependencies can be installed as external resources (<i>i.e.</i> , not through the package manager).	58%
<i>Usage of ENV</i>	Boolean	Environment variables must be supported for customizing the container.	83%
<i>Usage of ARG</i>	Boolean	Arguments must be supported for customizing the image build process.	75%
<i>Usage of LABEL</i>	Boolean	Labels must be used for documenting the Dockerfile.	50%
<i>Usage of EXPOSE</i>	Boolean	Network ports used must be documented.	67%
<i>Usage of CMD</i>	Boolean	The command to execute when starting the container must be specified through the CMD or	75%
<i>Usage of ENTRYPOINT</i>	Boolean	ENTRYPOINT command.	75%

TABLE I: Format of HLSes, with the type of each field, a description, and the percentage of survey participants who indicated the field as important.




Fig. 2: Steps performed to train T5 for generating Dockerfiles from specifications.

compose.yml file from arbitrary source code. Some other tools provide a more in-depth support for specific programming languages, based on the project context. For example, there are tools specific for R [6], Node.js [9], Ruby [8], and PHP [21]. Others support multiple languages, such as *generator-docker* [22].

To the best of our knowledge, no previous work introduced an approach that is able to generate complete Dockerfiles given high-level requirements.

III. DEEP LEARNING FOR GENERATING DOCKERFILES

We define a procedure to train a DL model for the generation of Dockerfiles from high-level specifications.

In Fig. 2 we report the workflow we used to train a DL-based model (T5) for the generation of Dockerfiles. We first define a structured high-level specification (HLS) for Dockerfiles. Then, we extract HLSes from existing Dockerfiles through an automated approach. Finally, we use such a dataset to train the T5 model for the generation of Dockerfiles (HLS \rightarrow Dockerfile). In the following, we describe in detail the steps to construct our model.

A. Dockerfile High-Level Specification

Natural language can be used as an effective tool for reporting the requirements of the source code. When it comes to Dockerfiles, however, the high-level requirements that can be expressed are much more limited. For the source code, a developer might want to specify, for example, constraints on the input parameters and conditions that lead to errors. On the other hand, for Dockerfiles, it boils down to a matter of what the developer wants installed in the container, plus a few more characteristics. Thus, to standardize the format of a Dockerfile requirements specification, written in natural language, the idea is to define a set of key-value requirements.

In a real-world application, it could be seen as a structured form where, for each field, the developer specifies the values to meet the requirements. Based on the commands available in Dockerfiles and on how Dockerfiles are generally structured (based on our experience), we distilled a structured format for high-level specifications (HLSes), which we report in Table I.

We did not include in the specification requirements related to particular Dockerfile commands (*i.e.*, ADD, COPY, HEALTHCHECK, MAINTAINER, ONBUILD, SHELL, STOPSIGNAL, USER, VOLUME, WORKDIR) for three reasons: First, they are related more to low-level details (*e.g.*, how the user of the Dockerfile should be set up); Second, in general, developers do not frequently use all of them [16], [23]; Third, some of them are deprecated (*i.e.*, MAINTAINER). Ideally, our approach automatically generates such low-level instructions when needed, *i.e.*, based on the instruction needed by the application that must be containerized. Also, developers might still easily tune them up on the generated Dockerfile, if they want to.

To validate the specification structure, we run a survey in which we asked 12 professional software developers with at least 2 years of experience with Docker whether they would want to specify each of the fields we hypothesized to be relevant in a HLS (binary yes/no question). We report the results in Table I. All the fields are important to at least 50% of the participants. The field which results to be relevant for the lower number of developers is *Usage of LABEL* (50% positive answers), while the most important ones are *Dependencies* and *Usage of ENV* (83% positive answers for both of them).

B. Inferring Requirements from Existing Dockerfiles

We need a large amount of associations HLS \rightarrow Dockerfile, to train a DL-based Dockerfile generation model. While Dockerfiles are largely available, this does not hold for the require-

ments behind them. The latter could be manually inferred, but such a process would be infeasible for a large-scale dataset. Thus, we defined an automated procedure for inferring the HLS behind an existing Dockerfile. In summary, the process works as follows. Given a list of operating systems (WL_{OS}) and stop words (WL_{stop}), we extract the operating system (step 1), the software dependencies (step 2), and, finally, the other fields required for HLS (step 3). In the following, we describe in detail each step of our methodology, and the procedure we used to define the two previously-mentioned lists of words.

1) *Step 1: Inferring the OS:* To infer the OS required by the developer, we only focus on the FROM instruction, which comes in the format FROM `<name>[:<tag>]`. Since `<name>` and `<tag>` are usually composed by one or more words, we extract them by splitting their content of image name and tag by the typically employed as separators, *i.e.*, `-` and `_`. We first check if any word in the WL_{OS} is present in the words extracted from the tag, and then in the name. If a OS-related keyword is present, we set *OS* with such a keyword, while we use the special keyword “any” otherwise, indicating that the developer did not have any requirement in terms of operating system. If more than a OS-related keyword is present, we consider the first match found. For example, for the instruction FROM `tomcat:9.0.20-jre8-alpine`, we extract the keywords (in the order): “9.0.20”, “jre8”, “alpine” (from the tag), and “tomcat” (from the image name). “alpine” is the first (and only, in this example) OS-related word. Therefore, we set *OS* to “alpine”. If the first OS-related keyword is found in the image name, we also add to it all the keywords containing only numbers from the tag name, which most likely refer to the version number. For example, for the instruction FROM `debian:10-slim`, we set *OS* to “debian10”.

2) *Step 2: Inferring Software Dependencies:* Software dependencies might be present in several Dockerfile instructions. The first instruction which might contain dependencies is the FROM instruction. To detect the (possible) dependencies explicitly expressed in such a field, we extract all the words from the image name (with the same procedure described in step 1), we remove the words in WL_{OS} (OSes) and WL_{stop} (stop words), and we exclude the non-alphabetic words.

All the remaining words are added to the list of software dependencies (*Dependencies*). In the previous example, FROM `tomcat:9.0.20-jre8-alpine`, the dependency extracted (and only word used as image name) is “tomcat”. At this point, we need to extract all the other dependencies, installed with package managers or with other procedure (*e.g.*, downloaded and installed). This task is far from trivial. Simply considering the packages installed through package managers (*e.g.*, `apt install`) might seem the natural step to achieve this goal. However, this is not an option: Most of the packages installed do not correspond to high-level requirements, but rather to low-level details about support libraries required to make some other software work. For example, in Fig. 1, the package `build-essential` is not a dependency of the software system, but rather a package incidentally needed (in this case, for building two actual dependencies, *i.e.*, `x265` and

`ffmpeg`). We want our DL model to automatically infer the need of such packages.

However, it is known that developers tend to give an explanation comment of what each Dockerfile instruction does. Thus, we can use those comments to extract only high-level requirements, which are reported by the developers themselves. To achieve this, we use the following heuristic. We first select all the comment lines that contain the word “install”. We tokenize each comment with the *spacy* Python library [24], and we detect the words that depend on such a keyword. We discard from the obtained list all the words in WL_{stop} , and we select the remaining ones as candidate dependencies for the specific comment line. For each comment line c_i with its candidate list of dependencies d_{c_i} , we process each RUN instruction between c and the next comment line (c_{i+1}) or blank line. We use *bashlex* [25] for parsing the bash script in each RUN instruction, and we split it in statements. Finally, we consider only the statements in which the most common package managers (*i.e.*, `apt`, `yum`, `apk`, `pip`, and `npm`) and commands for downloading files (*i.e.*, `curl` and `wget`) are present, and we extract their arguments. If there any argument (packages) matches a candidate requirement (from the comment), we add the candidate requirement in *Dependencies*. Note that we do not simply consider the words depending on the “install” word from the comments because some words might not be software dependencies (*e.g.*, in “install only ruby,” the word “only” should be ignored). Finally, we exclude duplicates, and we set the *Dependencies* with the dependencies extracted both from the FROM instruction and from the RUN instructions.

3) *Step 3: Parsing Additional Fields:* All the remaining fields are straightforward to be set. For the fields *Usage of ARG*, *Usage of CMD*, *Usage of ENTRYPOINT*, *Usage of ENV*, *Usage of EXPOSE*, and *Usage of LABEL*, we simply check if at least one of the respective instructions is present in the Dockerfile. The field *Package Manager* is defined by checking if any of the most commonly used package managers for Linux distributions (*i.e.*, `apt`, `apk`, and `yum`) are used in the RUN instructions. If this is the case, we set *Package Manager* to such a package manager (*e.g.*, *Package Manager* = `yum`). We also check for the coherence between the package manager and the OS (Linux distribution) detected in step 1: For example, if the OS is `ubuntu`, the package manager can not be `yum`. If this happens, or if no specific package manager is detected (*e.g.*, no package is installed), we set *Package Manager* to `any`. Finally, for the field *Download of External Dependencies*, we check if any RUN instruction contains one of the following: (i) a link in the context of a download-related instruction (*e.g.*, `wget` or `git clone`); (ii) the installation of a Python or JavaScript external library; (iii) the installation of an external package through the package managers (*e.g.*, `dpkg` for Debian/Ubuntu). If this happens, we set *Download of External Dependencies* to `true` (`false` otherwise).

4) *Defining OSes and Stop Words:* We use a systematic procedure for defining the two list required by the HLS parser, *i.e.*, WL_{OS} and WL_{stop} . We extracted all the FROM instructions

contained in the Dockerfiles from the collection by Eng *et al.* [16], that we later use to build our dataset. In total, we obtained 10,960,563 instructions. Then, we extract keywords from the image name and tag, like we do in step 1. Next, we discard the words that (i) contain only non-alphabetic characters (*e.g.*, version numbers), or (ii) have less than 3 characters (most likely stop words). We obtain a list of 46,070 unique words, along with the respective count of occurrences. We filter out all the words with less than 150 occurrences, thus obtaining 2,120 words, covering 97% of the total occurrences of all the words extracted. We manually analyzed and labeled each of them as “OS”, “stop word”, or “dependency”. For example, “centos” is marked as “OS” keyword, “baseimage” as “stop word”, while “python” as “dependency”. In the end, we obtained 920 stop words for WL_{stop} and 74 words referring to OSs for WL_{OS} .

C. Defining a Dataset of HLSes and Dockerfiles

We proceed with using our parser to build the dataset of HLSes and target Dockerfiles. We use the dataset built by Eng *et al.* [16], which is the largest (9.4M) and the latest dataset of Dockerfiles currently available in the literature. That dataset comes from the S version of World of Code (WoC) [26], covering a period of time ranging between 2013 and 2020. To the best of our knowledge, this is the most recent and large collection of Dockerfiles available in the literature. The dataset provides all the versions (snapshots) of all the Dockerfiles extracted from GitHub projects, along with the related commit ID. For building our dataset, we only select the latest commit for each repository. As a result, we obtain 3,010,141 Dockerfiles.

We filter out all the Dockerfiles that have no comments (required in step 2 of the parser), and also those that have more than one stage, as they are currently not supported by our parser. Next, we drop all the duplicated Dockerfiles (according to their *sha1* hash) and all the Dockerfiles that contain in the FROM instruction at least a keyword that we did not manually evaluate when defining the two lists WL_{OS} and WL_{stop} , to ensure that we do not unintentionally include stop words as dependencies. After applying these filters, we obtain a total of 769,385 Dockerfiles, on which we run our approach to extract the HLSes.

At this stage, we further exclude all the Dockerfiles for which our parser detected syntax errors in bash instructions and empty Dockerfile instructions (*e.g.*, COPY without arguments). As a result, we obtain a total of 670,982 pairs of Dockerfile, associated with the respective HLSes (in total, we found 121,030 unique HLSes). At this stage, each HLS is associated with one or more Dockerfiles since different Dockerfiles might result from the same requirements. We need, however, to select a single representative Dockerfile for each HLS that we can use for the training and the test of T5. To do this, we first select all HLS associated with two or more Dockerfiles. We found 41,820 of them. Given a set of Dockerfiles associated with the same HLS, we select the one containing the highest number of “typical” instructions for that cluster, to obtain the most typical Dockerfile for the given HLS. We tokenize each

instruction of all the Dockerfiles. Then, given two Dockerfiles A and B , we compute the Jaccard similarity between each pair of instructions of the same kind (*e.g.*, COPY instructions are compared only with COPY instructions), with the formula $J(W_{A_i}, W_{B_j}) = \frac{|W_{A_i} \cap W_{B_j}|}{|W_{A_i} \cup W_{B_j}|}$, where W_{A_i} and W_{B_j} are the words from instructions i and j of A and B , respectively. We choose as representative the Dockerfile with the highest mean similarity over all the instructions.

We further process the Dockerfiles to prepare them to the training procedure: First, given all the package installation instructions (*e.g.*, apt install), we sort the packages in lexicographic order, to avoid that different package orders confuse the DL approach. Second, we remove all the lines that contain only comments, to avoid that the model makes extra efforts in solving a sub-problem which is not in the scope of this paper.

From this, we extract two sub-datasets: D_{PT} for the pre-training, and D_{FT} for the fine-tuning of the T5 model. As for the former, we use all the Dockerfiles discarded while choosing the representative Dockerfiles for each HLS and all the Dockerfiles that are associated with a HLS for which the field *Dependencies* is empty (no software dependency needs to be installed). It is most likely that such Dockerfiles do have dependencies, but we were not able to find them because of the lack of comments in the format we expected (*e.g.*, the word “install” has not been used in comments). All the remaining pairs (HLS, Dockerfile) are placed in D_{FT} . In the end, we obtain 557,540 instances for D_{PT} and 113,442 instances for D_{FT} . As a requirement for the training of T5, we can only use Dockerfiles having no more than 1024 token, obtaining a total of 113,131 instances. Then, we divide the D_{FT} in training- ($D_{\text{FT-train}}$), evaluation- ($D_{\text{FT-eval}}$), and test- ($D_{\text{FT-test}}$) sets, by performing a typical 80%-10%-10% splitting [13], [14], obtaining 90,504, 11,313, and 11,314 instances, respectively. We use $D_{\text{FT-train}}$ for fine-tuning T5, $D_{\text{FT-eval}}$ for the hyper-parameter tuning, and $D_{\text{FT-test}}$ for our experiment (see Section IV).

D. Training T5 for Generating Dockerfiles

Raffel *et al.* [17] introduced T5 to support multitask learning in Natural Language Processing. Such a model re-frames NLP tasks in a unified text-to-text format in which the input and output of all tasks are always text strings. A T5 model is trained in two phases: (i) *pre-training*, in which the model is trained with a self-supervised objective that allows defining a shared knowledge-base useful for a large class of text-to-text tasks, and *fine-tuning*, which specializes the model on a specific downstream task (*e.g.*, learning the translation of sentences between different languages). As previously said, T5 already showed its effectiveness in code-related tasks [13], [27]–[32]. However, its application to the generation of Dockerfiles is novel and still unexplored. As done in previous work [13], [27], we use the smallest T5 version available (T5 small), which is composed of 60M parameters. Given a prediction provided by the model, the output token streams can be generated using several decoding strategies. We use *greedy*

decoding when generating an output sequence. In detail, such a strategy selects, at each time step t , the symbol having the highest probability of appearing in a specific position. We describe below both the pre-training and the fine-tuning procedure we applied for this task.

1) *Pre-Training Procedure*: The “general knowledge” [17] that we want to provide our model with is, in our case, a mixture of technical natural language (English) and technical language (Dockerfiles). We experiment with three pre-training variations: $T5_{NL}$, which only relies on natural language, $T5_{DF}$, which only relies on Dockerfiles, and $T5_{NL+DF}$, which relies on both. We test all such three variants and we pick the best after performing hyper-parameter tuning.

As for the first variant, $T5_{NL}$, we use the pre-trained checkpoint [33] released by Raffel *et al.* [17]. We do not perform any further pre-training for such a model.

Instead, we leverage the knowledge that has been already gained when pre-training the T5 model on the English text (C4 corpus [17]) for 1M steps.

As for the second variant, $T5_{DF}$, we adopt a classic *masked language model* task, *i.e.*, we randomly mask 15% of the tokens in a training instance, asking the model to predict them. We pre-train such a model on D_{PT} . Finally, as for the third variant, $T5_{NL+DF}$, we start from the $T5_{NL}$ model and we further pre-train it for 500k steps on D_{PT} , using the same procedure used for pre-training $T5_{DF}$. Finally, we created a new *SentencePiece* model [34] for tokenizing natural language text. We trained it on D_{PT} . For both the models for which we performed additional pre-training steps ($T5_{DF}$ and $T5_{NL+DF}$), we used a 2x2 TPU topology (8 cores) from Google Colab with a batch size of 16 and a sequence length of 512 tokens for the input and 1,250 for the output. As a learning rate, we use the Inverse Square Root with the default configuration [17]. For the pre-training phase, we use the default parameters defined for the T5 model [17].

2) *Hyper-parameter Tuning*: We test four learning rate strategies, *i.e.*, constant learning rate (C-LR), slanted triangular learning rate (ST-LR), inverse square learning rate (ISQ-LR), and polynomial learning rate (PD-LR). We report in Table II the parameters we use for each of them.

Given the three pre-trained models, $T5_{NL}$, $T5_{DF}$, and $T5_{NL+DF}$, we fine-tune them on $D_{FT-eval}$ (100k steps, batch size of 32, input sequence length of 512 tokens, output sequence length of 1024 tokens), leading to 12 different models (3 models \times 4 strategies). Then, to assess their performance, we compute the BLEU-4 [35] metric between the generated Dockerfiles and the target ones. Such a metric has been used in previous work for other coding tasks [13], [14], [36], [37] and it ranges between 0 (completely different) and 1 (identical). We report in Table III the results achieved by the 12 models in terms of BLEU-4. The best results are achieved with $T5_{DF}$ with the ISQ-LR strategy and $T5_{NL+DF}$ with the ST-LR strategy (17.20% BLEU-4 for both). In the end, we select the latter since $T5_{NL+DF}$ achieves better results also for the other strategies.

TABLE II: Configurations for the experimented learning rates

Strategy	Parameters	Strategy	Parameters
C-LR	$LR = 0.001$	ISQ-LR	$LR_s = 0.01$
ST-LR	$LR_s = 0.001$		$W = 10,000$
	$LR_{max} = 0.01$	PD-LR	$LR_s = 0.01$
	$Ratio = 32$		$LR_e = 0.001$
	$Cut = 0.1$		$Pow = 0.5$

TABLE III: T5 hyper-parameter tuning results (BLEU-4).

Experiment	C-LR	ST-LR	ISQ-LR	PD-LR
$T5_{NL}$	13.80%	13.70%	5.50%	14.50%
$T5_{DF}$	16.90%	5.50%	17.20%	15.90%
$T5_{NL+DF}$	16.60%	17.20%	16.50%	17.10%

3) *Fine-tuning*: We fine-tune the best pre-trained model ($T5_{NL+DF}$) with the best learning rate strategy (ST-LR) on $D_{FT-train}$. We use early stopping to avoid overfitting [29], [38]: We save a checkpoint every 10k steps and compute the BLEU-4 score on the evaluation set every 100k steps. When the 100k steps do not lead to an improvement, we stop the training procedure, and we keep the last model.

IV. EMPIRICAL STUDY DESIGN

The *goal* of our study is to understand to what extent T5 is effective in generating Dockerfiles. Our study is steered from the following research questions:

RQ₁: *To what extent is T5 able to generate Dockerfiles meeting the input natural language specification?* We evaluate the effectiveness of T5 in generating Dockerfiles that meet the requirements.

RQ₂: *To what extent are the Dockerfiles generated by T5 similar to the original ones written by developers?* With this second RQ, we aim at understanding if T5 generates Dockerfiles similar to the targets ones.

RQ₃: *To what extent are the Docker Images built from the Dockerfiles generated by T5 similar to the original ones built from the Dockerfiles written by developers?* Two Docker images can be equal and come from completely different Dockerfiles. Therefore, with this last RQ, we try to understand whether the images built from the generated Dockerfiles are similar to the original ones.

A. Baseline Techniques

We use as baseline techniques two Information Retrieval (IR)-based approaches. The first one is IR_{ES} , and it is based on the state-of-the-practice for implementing IR approaches and search engines, *i.e.*, Elasticsearch [39]. Given a collection of documents (D) and a query (q), Elasticsearch first assigns a score to each document in D according to q and then it sorts them. The score is computed with Okapi BM25 [40]. We add into an Elasticsearch instance all the instances in $D_{FT-train}$. Specifically, for each instance, we define a document that contains both the HLS and the associated Dockerfile. Given a new HLS for which we want to get a candidate Dockerfile, we perform a *boolean query* composed of all the fields of the

HLS in OR clause (*i.e.*, `should`, in Elasticsearch). We report an example of the query in our replication package [41].

The second baseline we consider is IR_{ST} , and it is based on the *SentenceTransformers* framework [42]. Such a framework allows to train embeddings for several data types (including text) so that they can be represented as numeric vectors. First, we use $D_{FT-train}$ to train the model for computing the embeddings (*bert-base-uncased*). Then, we compute the embeddings for each HLS as $e(d_{Spec})$ for each HLS in $D_{FT-train}$, and we store their associations with the respective Dockerfiles ($e(d_{Spec}) \rightarrow d_{Dockerfile}$). Given a new HLS, t_{Spec} , we compute its embeddings ($e(t_{Spec})$) and, then, the cosine similarity between $e(t_{Spec})$ and each $e(d_{Spec})$. Finally, we return the $d_{Dockerfile}$ for which the aforementioned similarity is maximum.

B. Context Selection

The *context* of our study is composed of (i) a set of associations $HLS \rightarrow Dockerfile$, (ii) the Dockerfiles generated/retrieved by T5 and the two baseline techniques, and (iii) the source code of the software projects for which the Dockerfiles need to be built, for building the images and thus answering RQ_3 .

As for the first object, we used the $D_{FT-test}$ dataset. To obtain the second object, we ran T5, IR_{ES} , and IR_{ST} on the HLSes from $D_{FT-test}$. As a result, we obtained three sets of generated/retrieved Dockerfiles, *i.e.*, DF_{T5} , DF_{ES} , DF_{ST} . Finally, to obtain the third object, for each Dockerfile in $D_{FT-test}$, we consider the original entry in the dataset by Eng *et al.* [16] and we recover the project from which it was extracted and the commit for that specific snapshot. We cloned all the repositories corresponding to the test instances, discarding the ones for which the source commit or repository was no longer available. As a result, we obtained a total of 3,909 repositories, corresponding to 4,059 instances of $D_{FT-test}$. Since building Dockerfiles requires a large amount of time, we did this for a representative sample of 500 Dockerfiles from $D_{FT-test}$ (4.28% margin of error, 95% confidence level). For each instance in $D_{FT-test}$, we tried to build the original Dockerfile: If the build failed, we discarded the instance, while, if it succeeded, we kept it, until we collected 500 instances.

C. Experimental Procedure

To answer RQ_1 , we compare the HLS related to the Dockerfiles returned by the three approaches we consider with the one given as input. As for the two baselines, we already have an associated HLS for each returned Dockerfile (*i.e.*, the one from $D_{FT-train}$). This, however, is not true for T5 since it generates Dockerfiles from scratch. In this case, we use the same process described in Section III-B on the generated Dockerfiles for all the fields, except for *Dependencies*. In this case, we check if requirements in the HLS given as input are met in the generated Dockerfile. The reason is that we trained T5 not to generate comments: Since our procedure for extracting *Dependencies* strongly relies on comments, we can not directly use it on the generated Dockerfiles. We ignore the

Dockerfiles generated by T5 for which the parser we defined is not able to infer all the fields in the HLS (1,337 instances, *i.e.*, $\sim 12\%$ of the total). We consider, instead, all of them for the two baselines, for which this cannot happen since, as mentioned, we do not infer the HLS. Finally, we measure the similarity between the target HLSes and the obtained ones. To achieve this, we assign a *score* for each field of the HLSes, compared to the respective instance in the target HLS, which is computed by assigning 1 point if the field is equal, and 0 otherwise. The only exception is the *Dependencies* field: Since this is a collection of elements, in this case we compute the score by computing the percentage of elements in the target HLS that are present also in the obtained one (*i.e.*, the recall). We compute and report the mean score obtained for each field of the HLSes.

To answer RQ_2 , we compare the sets DF_{T5} , DF_{ES} , DF_{ST} with the respective target Dockerfiles. Simply computing textual similarity for Dockerfiles is not sufficient since, in many cases, instructions can be swapped without affecting the final result. Therefore, we rely on the AST representation of the Dockerfiles. To do this, we use the *binnacle* tool by Henkel *et al.* [43]. *Binnacle* extracts ASTs at three different abstraction levels. We use the *phase-2* representation. We do not use the *phase-3* abstraction level since it abstracts the bash commands. For example, both the `apt-get install` and `pacman -S` instructions get replaced with a generic *package install*. While such an abstraction is useful, the tool does not support all the possible bash commands, thus causing loss of information, in some cases. We report an example of a parsed AST in our replication package [41]. Given the ASTs of two Dockerfiles, we compute the *edit distance* between them, *i.e.*, the number of modifications needed to transform one into the other, using the Zhang-Shasha algorithm [44]. We normalize the obtained edit distance by dividing it by the sum of the sizes of the two trees we are comparing. We discard all the instances for which the *binnacle* tool is not able to extract the AST. We obtain 9,963 for T5, 11,311 for IR_{ES} , and 11,129 for IR_{ST} . We run the Mann–Whitney U test [45] to compare T5 with the baselines in terms of normalized edit distance. The null hypothesis is that there is no difference between the edit distance obtained using T5 and the one obtained using an IR-based approach. We correct for multiple comparisons using the Benjamini-Hochberg procedure [46]. Finally, we compute the effect size using Cliff’s Delta [47], which is negligible for all of the performed comparison (*i.e.*, T5 compared with the two baselines). This means that their difference is negligible, even if it is statistically significant.

To answer RQ_3 , we compare the Docker images built from the resulting Dockerfiles provided by the three techniques and the one built from the target Dockerfile. We consider the GitHub projects we cloned for the sample of 3,909 instances and, for each of them, we replace the original Dockerfile with the one generated/retrieved by the three approaches, one at a time, and we try to build it. For each instance, we memorized (i) if the build succeeded, (ii) if the the original and the obtained images are equal, and, if not, (iii) to what extent the

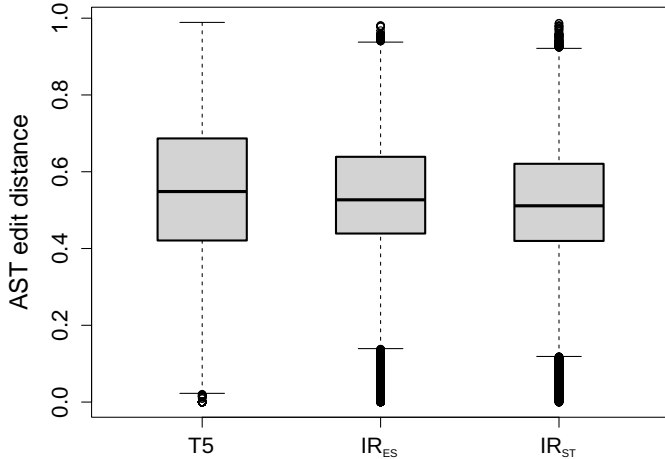


Fig. 3: Boxplots of the normalized AST edit distance (RQ₂).

latter provides what is also present in the former. As for the second measurement, we rely on the image digest: We say that two images are equal if their digest are equal. While this is not true by design, we can safely assume that, in our context, the risk of obtaining different images with the same digest is negligible. As for the third measurement, we compute the digest of each build layer (*i.e.*, one for each Dockerfile instruction), and we compute the percentage of layer digests of the image resulting from the original Dockerfile that also appear in the image built from the generated/retrieved Dockerfile. Note that if this measure is 100%, in this context, it means that the generated/retrieved Dockerfile is able to provide everything that the original image already provided. Still, it is possible that it contains additional layers not present in the original image.

D. Data Availability

We provide a replication package [41] containing the tool to extract high-level requirements from Dockerfiles, the trained DL models, along with the datasets and the code for replicating the baseline techniques and the conducted experiment.

V. EMPIRICAL STUDY RESULTS

In this section, we report the results of our study and, thus, the answers to our research questions.

A. RQ₁: Adherence to the High-Level Specification

In Table IV, we report the results of the comparison with input HLS fields. First, it can be noticed that IR_{ES} is the best-performing baseline, since it always achieves better or equal results than IR_{ST}. Therefore, from now on, we only discuss the comparison between T5 and such a baseline in this RQ. T5 is generally able to better meet the requirement in terms of *OS* (+7.6pp), while it achieves slightly worse results in terms of *Package Manager* (-1.9pp), *Dependencies* (-1.2pp), and *Download of External Dependencies* (-2.6pp). The last two are probably the most critical and hard-to-meet requirements

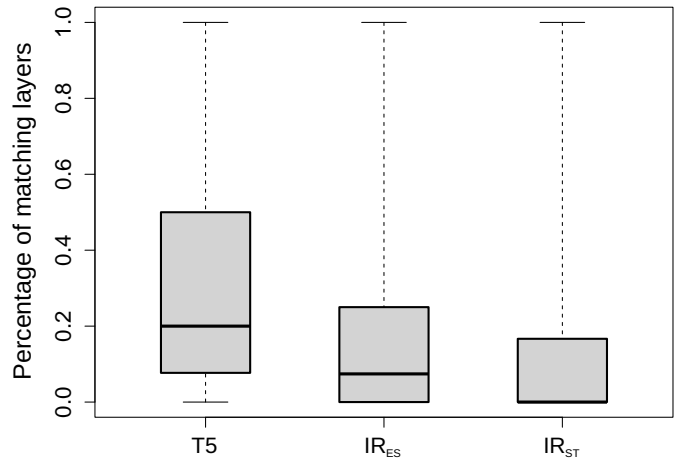


Fig. 4: Boxplots of the percentage of matching layers (RQ₃).

since they also interact with each other, and we can observe that both the approaches generally achieve good results. As for the other requirements, we observe that T5 performs better on *Usage of ENV*, *Usage of ARG*, and *Usage of LABEL*, while IR_{ES} achieves better results in terms of *Usage of EXPOSE*, *Usage of CMD*, and *Usage of ENTRYPOINT*. In summary, we can conclude that (i) there is no clear winner between the two approaches, and (ii) both the approaches generally return Dockerfiles that meet most of the requirements.

In summary, T5 and IR_{ES} perform very similarly in terms of adherence to the input requirements: There is no clear winner as for this aspect.

B. RQ₂: Dockerfile Similarity

We report the adjusted boxplots [48] for the normalized *AST edit distance* in Fig. 3. The boxplot shows the distribution of the edit distance between the generated Dockerfile and the original Dockerfiles for each instance of the test set. The higher the distance, the lower the similarity. Also in this case, it seems that there is no clear winner: T5 has higher variance, thus being able to generate both better and worse Dockerfile compared to the two baselines. The mean edit distance is 0.55 ($\sigma = 0.19$) for T5, 0.53 ($\sigma = 0.18$) for IR_{ES}, and 0.51 ($\sigma = 0.19$) for IR_{ST}. This means that the two IR-based baselines perform slightly better than T5. The difference is significant according to the Mann–Whitney U tests we performed for comparing T5 with the IR_{ES} and IR_{ST} (adjusted *p*-value lower than 0.001 for both). The Cliff’s Delta between T5 and IR_{ES} is 0.06, and 0.11 between T5 and IR_{ST}. Thus, the difference is *negligible* in both cases. We took a closer look at the cases in which the generated/retrieved Dockerfile was perfectly equal to the original one (*i.e.*, edit distance 0). We have 93 of such cases for T5, while only 18 and 11 for IR_{ES} and IR_{ST}, respectively. In terms of AST size, the perfect matches are rather small (10.1, $\sigma = 9.4$, with a maximum of 58) compared to the average size (133.7, $\sigma = 132.5$) for T5, while it is remarkably higher for the two baselines (39.5, $\sigma =$

TABLE IV: Adherence score between the input and the generated HLS reported for each field.

Approach	OS	Pkg Man.	Dep.	ENV	ARG	LABEL	EXPOSE	CMD	ENTRYPOINT	Down. Ext. Dep.
T5	0.998	0.981	0.865	0.892	0.987	0.999	0.798	0.743	0.843	0.816
IR _{ES}	0.922	1.000	0.877	0.812	0.884	0.872	0.829	0.826	0.851	0.842
IR _{ST}	0.880	1.000	0.761	0.518	0.168	0.165	0.373	0.453	0.260	0.448

38.2 for IR_{ES}, and 41.8, $\sigma = 44.8$ for IR_{ST}). Such a result suggests that T5 works well when small Dockerfile need to be generated, while it struggles with bigger ones. To confirm such a conjecture, we computed the correlation (Spearman ρ) between the AST size of the target Dockerfile for T5, and we found that it is significant and high ($\rho = 0.67$), much more than the two baselines ($\rho = -0.12$ for IR_{ES}, and not significantly different from 0 for IR_{ST}).

In summary, all the approaches returns Dockerfiles quite different from the target ones. T5 works well with small Dockerfiles, but not with bigger ones.

C. RQ₃: Docker Images Similarity

T5 achieves a build success rate of 34% (170/500 correctly built images), outperforming the IR_{ES} (23%, 166/500 images) and IR_{ST} (32%, 156/500 images). Comparing the *digest* of the built images (*i.e.*, hash value) with the source image (*i.e.*, the one built from the test instance), we obtain remarkably better results for T5: We have 11.7% of matches (20/170 instances), while the two baselines have no matching digest for their images. This result is confirmed also in Fig. 4, which depicts the distribution of the percentage of matching layers (adjusted boxplots [48]). The mean percentage of matching layers is 32.4% for T5 ($\sigma = 0.32$), 17.5% for IR_{ES} ($\sigma = 0.26$), and 10.9% for IR_{ST} ($\sigma = 0.19$). The obtained results complement the ones presented in RQ₂. Not only T5 works better when it needs to generate small Dockerfiles, but it also works better than the two baselines on bigger ones, up to a certain point; after that, it is not able to generate good instructions, thus the limited layer match, in absolute terms.

In summary, T5 achieves the best results compared to the two baselines in terms of build success, percentage of perfectly matching images, and percentage of matching layers.

VI. DISCUSSION

Generating Dockerfiles from high-level specifications is a challenging task. In this study, we perform a first attempt to solve this problem using Deep Learning (T5, specifically). Considering the overall results, there is no clear evidence that T5 is better than using IR-based techniques, in practice. Given the lower effort in setting-up a IR-based technique (which is trivial in the case of IR_{ES}, for example), at a superficial level, we can conclude that, at the moment, this would be the best option for practitioners.

However, we analyzed the results more in-depth to try to understand what went wrong, and why T5 does not work well for this task, despite it works very well for other coding tasks [14]. First, we observed that T5 generates Dockerfiles with a much lower number of tokens compared to the two baseline approaches (36.80 vs. 135.70 for IR_{ES} and 107.10 for IR_{ST}).

```

1 FROM golang:1.9.4-stretch
2
3 RUN apt-get update -y && apt-get upgrade -y
4
5 RUN git clone https://github.com/edenhill/librdkafka.git && \
6   cd librdkafka && \
7   ./configure --prefix /usr && \
8   make && \
9   make install && \
10  cd ..
11
12 WORKDIR /go/src/consumerpg
13 COPY . .
14
15 RUN go-wrapper download
16 RUN go-wrapper install
17
18 CMD ["go-wrapper", "run"]

```

(a) Target Dockerfile.

```

1 FROM golang:1.9.4-stretch
2
3 RUN apt-get update && apt-get install -y git
4
5 RUN git clone https://github.com/edenhill/librdkafka.git && cd
6   librdkafka && ./configure && make && make install
7
8 RUN go get github.com/confluentinc/confluent-kafka-go
9
10 WORKDIR /go/src/github.com/confluentinc/confluent-kafka-go
11 CMD ["go", "run", "main.go"]

```

(b) Dockerfile generated by T5

Fig. 5: Example of a generated incomplete Dockerfile

This explains why T5 works well for smaller Dockerfiles and gradually less well for bigger ones (RQ₂), and also why T5 achieves a good percentage of matching layers even if the Dockerfile similarity is low (RQ₃).

We manually analyzed some Dockerfiles generated with T5. We found that, in some instances, the Dockerfiles abruptly interrupt in the middle of the last instruction. While the remainder of the generated Dockerfiles is correct, the last instruction often contains issues. An example is provided in Fig. 5b, with the target Dockerfile (a) and the one generated by T5 (b). The prediction is remarkably good, until, in the last line, T5 stops the generation at a certain point. A pattern we observed is that interrupted Dockerfiles do not end with the token we used for indicating the new line ($\langle n1 \rangle$), while the ones in the training always end with such a token by design. We counted a total of 6,786 instances of such a kind ($\sim 60\%$ of the cases).

If we consider only the instances that terminate with the newline token, the results of T5 become better than the two IR baselines in all the aspects we considered: For RQ₁, T5 achieves better results than both the baselines for all the fields, except for *Package Manager*; for RQ₂, the edit distance becomes significantly lower, *i.e.*, 0.46, compared to 0.51 of the best baseline for such a sub-sample; For RQ₃, we obtain results in line with the previously presented ones. We tried to address this issue by replacing the greedy decoding strategy with a sampling strategy, with different values for the *temperature*

hyper-parameter of the *softmax* function. A high *temperature* allows to increase the chances of picking tokens with lower likelihood, while a lower *temperature* does the exact opposite, so that, when the temperature is close to 0, such a decoding strategy behaves like a greedy decoding strategy.

We tested temperature values between 0.7 and 1.0, with a step of 0.1. We observe that increasing the *temperature* allows to reduce the number of incomplete Dockerfiles: With a *temperature* of 0.7, we obtain a total of 1,648 incomplete generations (14.6%) which decrease to 1,283 (11.3%) with a *temperature* of 1. The average number of tokens contained in the Dockerfiles varies between 82 (*temperature* = 1.0) and 84 (*temperature* = 0.7), while with greedy decoding we have 37, on average. In the end, however, the Dockerfiles generated with this strategy achieve generally worse results in terms of (i) number perfect predictions (\sim 28%, with the best *temperature*, *i.e.*, 0.8), (ii) number of perfectly matching images (\sim 80%, best *temperature* = 1.0) and layers (\sim 20%, best *temperature* = 0.9).

This analysis shows that, while T5 learned how to generate Dockerfiles, to some extent, it does not have enough knowledge to generate complete Dockerfiles well: It either partially generates good parts of Dockerfiles or generates complete less-good Dockerfiles. There are two possible explanations for this phenomenon, which also represent open problems for this specific task:

- **A larger dataset needs to be built.** Addressing this problem is only apparently easy. We considered in our study the largest collection of Dockerfiles available in the literature [16], which includes *all* the Dockerfiles from open-source projects produced up to 2020. While such a dataset can be updated with the last two years of activities in GitHub, it is unlikely that the size of our dataset would drastically increase as a result. Indeed, we consider a single Dockerfile for each unique HLS, *i.e.*, we would not have new instances for the HLSes already covered. We rely on comments by the authors to extract some requirements (specifically, the *Dependencies* field). Some Dockerfiles, however, do not have explicit indications of such requirements and, therefore, we miss them in the inferred specifications. New and more precise ways of extracting high-level requirements from Dockerfiles are needed. Also, a promising direction would consist in the definition of techniques for data augmentation in this context, *e.g.*, by blending existing Dockerfiles to provide uncovered combinations of HLSes.
- **A different training stop criterion needs to be defined.** In this study, we used the same procedure previously used for coding tasks [14], for which the problem of abrupt interruption in the inference has not been observed. It is possible that the stop criterion and the metric (*BLEU*) used are not the right ones in this context. As for the latter, it might be worth exploring different distance measures. While the AST distance is not a viable option for performance reasons, other metrics that do not consider the order in the instructions might be more useful.

VII. THREATS TO VALIDITY

Threats to Construct Validity concern the correct operationalization of the concepts being studied. First, the inference method we used for extracting requirements from Dockerfiles works on a series of assumptions (*e.g.*, the presence of comments) that might not always be completely satisfied in practice. To mitigate this crucial threat, we carefully tested our parser and manually checked examples until we were satisfied with the procedure used to achieve this goal. In total, we were able to infer 113,442 unique HLSes, with 31,990 unique combination of dependencies, which gives us confidence on the fact that our parser works as intended for most of the Dockerfiles considered. Also, we made sure to exploit all the instances (also the ones without comments) in the training procedure (*i.e.*, by using them for pre-training). The choice of the fields that compose our HLS could exclude requirements that developers might be interested in specifying (such as *USER*). As we explained in Section III, we exclude only the requirements derived from Dockerfile instructions that do not appear very often. The best candidate Dockerfile we selected for each HLS (Section III) might not be representative of the group of Dockerfiles. In our methodology, we relied on Jaccard similarity to mitigate this risk.

Threats to Internal Validity concern factors internal to our study that might have influenced our findings. The percentage of marching layers (RQ_3) might not accurately capture the structural similarity between two Docker images: If a layer is different, many or even all the subsequent layers will be different. However, to the best of our knowledge, the only alternative is to perform a diff on Docker containers [49], which, however, only works at the level of installed packages. Considering the layers allows knowing from which point the two images started to differ, we believe this is the best choice.

Threats to External Validity concern the generalizability of our findings. We relied on the largest collection of Dockerfiles available in the literature [16]. More Dockerfiles might have been created between 2020 and now. However, we believe that the considered dataset allows us to provide reasonably generalizable.

VIII. CONCLUSION

We evaluated the effectiveness of Deep Learning (and, specifically, T5) for the automatic generation of Dockerfiles. The results show that, while T5 works very well on small Dockerfile, it struggles with larger ones. After having deeply analyzed this phenomenon, we identified two possible issues that must be addressed before deploying a working DL-based solution for this task. First, it is necessary to *build a larger dataset*, which is not easy, given that we used the largest collection of open-source Dockerfiles available. Second, it is necessary to *devise a different stopping criterion* for fine-tuning T5 since the one typically used for coding tasks (based on BLEU-4) likely causes an early stop, which does not allow the model to properly complete the learning process.

REFERENCES

- [1] J. Humble and D. Farley, "Continuous delivery: reliable software releases through build," *Test, and deployment automation. Pearson Education*, vol. 1, 2010.
- [2] StackOverflow, "2022 developer survey," <https://survey.stackoverflow.co/2022/#section-most-loved-dreaded-and-wanted-other-tools>, [Online; accessed 7-Jul-2022].
- [3] D. Reis, B. Piedade, F. F. Correia, J. P. Dias, and A. Aguiar, "Developing docker and docker-compose specifications: A developers' survey," *IEEE Access*, vol. 10, pp. 2318–2329, 2021.
- [4] "Broken by default: why you should avoid most dockerfile examples," <https://pythonspeed.com/articles/dockerizing-python-is-hard/>, [Online; accessed 31-Aug-2022].
- [5] E. Horton and C. Parnin, "Dockerizeme: Automatic inference of environment dependencies for python code snippets," in *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*. IEEE, 2019, pp. 328–338.
- [6] D. Nüst and M. Hinz, "containerit: Generating dockerfiles for reproducible research with r," *Journal of Open Source Software*, vol. 4, no. 40, p. 1603, 2019.
- [7] "starter: Helping you get started with containerized apps," <https://github.com/cloud66-oss/starter/commit/3b10a910a5cff8cfe02c628d08ad373018cac888>, 2015, [Online; accessed 31-Aug-2022].
- [8] "boxing: The zero-configuration dockerfile generator for ruby," <https://github.com/elct9620/boxing/commit/4e3a200c2420bf5a5108f3061a27c0c1903cee6e>, 2021, [Online; accessed 31-Aug-2022].
- [9] "dockerfile-generator: Dockerfile generator for node.js," <https://github.com/tudvari/dockerfile-generator/commit/fd4582a7a1bb512b518c9a130e9552378f4f601b>, 2018, [Online; accessed 31-Aug-2022].
- [10] H. Ye, J. Zhou, W. Chen, J. Zhu, G. Wu, and J. Wei, "Dockergen: A knowledge graph based approach for software containerization," in *2021 IEEE 45th Annual Computers, Software, and Applications Conference (COMPSAC)*. IEEE, 2021, pp. 986–991.
- [11] "Github copilot: Your ai pair programmer," <https://github.com/features/copilot>, [Online; accessed 31-Aug-2022].
- [12] M. Tufano, C. Watson, G. Bavota, M. D. Penta, M. White, and D. Poshyvanyk, "An empirical study on learning bug-fixing patches in the wild via neural machine translation," *ACM Transactions on Software Engineering and Methodology (TOSEM)*, vol. 28, no. 4, pp. 1–29, 2019.
- [13] A. Mastropaolo, S. Scalabrino, N. Cooper, D. N. Palacio, D. Poshyvanyk, R. Oliveto, and G. Bavota, "Studying the usage of text-to-text transfer transformer to support code-related tasks," in *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)*. IEEE, 2021, pp. 336–347.
- [14] A. Mastropaolo, N. Cooper, D. N. Palacio, S. Scalabrino, D. Poshyvanyk, R. Oliveto, and G. Bavota, "Using transfer learning for code-related tasks," *IEEE Transactions on Software Engineering*, 2022.
- [15] C. Watson, N. Cooper, D. N. Palacio, K. Moran, and D. Poshyvanyk, "A systematic literature review on the use of deep learning in software engineering research," *ACM Transactions on Software Engineering and Methodology (TOSEM)*, vol. 31, no. 2, pp. 1–58, 2022.
- [16] K. Eng and A. Hindle, "Revisiting dockerfiles in open source software over time," in *2021 IEEE/ACM 18th International Conference on Mining Software Repositories (MSR)*. IEEE, 2021, pp. 449–459.
- [17] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, P. J. Liu *et al.*, "Exploring the limits of transfer learning with a unified text-to-text transformer," *J. Mach. Learn. Res.*, vol. 21, no. 140, pp. 1–67, 2020.
- [18] "Dockerfile reference," <https://docs.docker.com/engine/reference/builder/>, [Online; accessed 31-Aug-2022].
- [19] K. Hanayama, S. Matsumoto, and S. Kusumoto, "Humpback: Code completion system for dockerfiles based on language models," in *Proc. Workshop on Natural Language Processing Advancements for Software Engineering*, 2020, pp. 1–4.
- [20] Y. Zhang, Y. Zhang, X. Mao, Y. Wu, B. Lin, and S. Wang, "Recommending base image for docker containers based on deep configuration comprehension," in *2022 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*. IEEE, 2022, pp. 449–453.
- [21] "phpdocker.io: Phpdocker.io website and environment (php and docker based) generator," <https://github.com/phpdocker-io/phpdocker.io/commit/4ea67367e969c23a7c3cdd42c749b2bd53e7b3c3>, 2016, [Online; accessed 31-Aug-2022].
- [22] "Yeoman generator for docker," <https://github.com/microsoft/generator-docker/commit/9bb74be96d3d9586ee86b8caec24241528c0bf33>, 2015, [Online; accessed 31-Aug-2022].
- [23] J. Cito, G. Schermann, J. E. Wittern, P. Leitner, S. Zumberi, and H. C. Gall, "An empirical analysis of the docker container ecosystem on github," in *2017 IEEE/ACM 14th International Conference on Mining Software Repositories (MSR)*. IEEE, 2017, pp. 323–333.
- [24] "spacy: Industrial-strength natural language processing (nlp) in python," <https://github.com/explosion/spaCy/commit/d583626a826c00dfba55f42dc7911d1a4b0b7032>, 2017, [Online; accessed 31-Aug-2022].
- [25] "bashlex: Python parser for bash," <https://github.com/idank/bashlex/commit/58fe9280c431491af0a8c4e5722c638072bab68e>, 2014, [Online; accessed 31-Aug-2022].
- [26] Y. Ma, C. Bogart, S. Amreen, R. Zaretski, and A. Mockus, "World of code: an infrastructure for mining the universe of open source vcs data," in *2019 IEEE/ACM 16th International Conference on Mining Software Repositories (MSR)*. IEEE, 2019, pp. 143–154.
- [27] A. Mastropaolo, E. Aghajani, L. Pascarella, and G. Bavota, "An empirical study on code comment completion," in *2021 IEEE International Conference on Software Maintenance and Evolution (ICSM)*. IEEE, 2021, pp. 159–170.
- [28] A. Mastropaolo, L. Pascarella, and G. Bavota, "Using deep learning to generate complete log statements," in *Proceedings of the 44th International Conference on Software Engineering*, ser. ICSE '22. New York, NY, USA: Association for Computing Machinery, 2022, pp. 2279–2290. [Online]. Available: <https://doi.org/10.1145/3510003.3511561>
- [29] R. Tufano, S. Masiero, A. Mastropaolo, L. Pascarella, D. Poshyvanyk, and G. Bavota, "Using pre-trained models to boost code review automation," in *44th IEEE/ACM 44th International Conference on Software Engineering, ICSE 2022, Pittsburgh, PA, USA, May 25-27, 2022*, 2022, pp. 2291–2302. [Online]. Available: <https://doi.org/10.1145/3510003.3510621>
- [30] L. Li, L. Yang, H. Jiang, J. Yan, T. Luo, Z. Hua, G. Liang, and C. Zuo, "Auger: Automatically generating review comments with pre-training models," *arXiv preprint arXiv:2208.08014*, 2022.
- [31] J. Zhang, S. Panthaplackel, P. Nie, J. J. Li, and M. Gligoric, "Coditt5: Pretraining for source code and natural language editing," *arXiv preprint arXiv:2208.05446*, 2022.
- [32] M. Ciniselli, N. Cooper, L. Pascarella, A. Mastropaolo, E. Aghajani, D. Poshyvanyk, M. Di Penta, and G. Bavota, "An empirical study on the usage of transformer models for code completion," *IEEE Transactions on Software Engineering*, 2021.
- [33] "T5 pre-trained public checkpoint," [gs://t5-data/pretrained_models/small](https://huggingface.co/google/t5-data/pretrained_models/small).
- [34] T. Kudo and J. Richardson, "Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing," *arXiv preprint arXiv:1808.06226*, 2018.
- [35] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu, "Bleu: a method for automatic evaluation of machine translation," in *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, 2002, pp. 311–318.
- [36] Z. Liu, X. Xia, A. E. Hassan, D. Lo, Z. Xing, and X. Wang, "Neural-machine-translation-based commit message generation: how far are we?" in *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering*, 2018, pp. 373–384.
- [37] A. LeClair, S. Jiang, and C. McMillan, "A neural model for generating natural language summaries of program subroutines," in *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*. IEEE, 2019, pp. 795–806.
- [38] Y. Yao, L. Rosasco, and A. Caponnetto, "On early stopping in gradient descent learning," *Constructive Approximation*, vol. 26, no. 2, pp. 289–315, 2007.
- [39] "Elasticsearch: The official distributed search and analytics engine," <https://www.elastic.co/elasticsearch/>, [Online; accessed 31-Aug-2022].
- [40] "Practical bm25 - part 2: The bm25 algorithm and its variables," <https://www.elastic.co/blog/>

practical-bm25-part-2-the-bm25-algorithm-and-its-variables, [Online; accessed 31-Aug-2022].

- [41] A. Authors, *Replication package*, TBD, <https://figshare.com/s/217ac28f442c0be36857>.
- [42] N. Reimers and I. Gurevych, "Sentence-bert: Sentence embeddings using siamese bert-networks," in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 11 2019. [Online]. Available: <https://arxiv.org/abs/1908.10084>
- [43] J. Henkel, C. Bird, S. K. Lahiri, and T. Reps, "Learning from, understanding, and supporting devops artifacts for docker," in *2020 IEEE/ACM 42nd International Conference on Software Engineering (ICSE)*. IEEE, 2020, pp. 38–49.
- [44] K. Zhang and D. Shasha, "Simple fast algorithms for the editing distance between trees and related problems," *SIAM journal on computing*, vol. 18, no. 6, pp. 1245–1262, 1989.
- [45] F. Wilcoxon, "Individual comparisons by ranking methods," *Biometrics Bulletin*, vol. 1, no. 6, pp. 80–83, 1945. [Online]. Available: <http://www.jstor.org/stable/3001968>
- [46] Y. Benjamini and Y. Hochberg, "Controlling the false discovery rate: a practical and powerful approach to multiple testing," *Journal of the Royal statistical society: series B (Methodological)*, vol. 57, no. 1, pp. 289–300, 1995.
- [47] N. Cliff, "Dominance statistics: Ordinal analyses to answer ordinal questions," *Psychological bulletin*, vol. 114, no. 3, p. 494, 1993.
- [48] M. Hubert and E. Vandervieren, "An adjusted boxplot for skewed distributions," *Computational statistics & data analysis*, vol. 52, no. 12, pp. 5186–5201, 2008.
- [49] "container-diff: Diff your docker containers," <https://github.com/GoogleContainerTools/container-diff/commit/0f743be182c3f38c5343d50676b7ef4048102921>, 2017, [Online; accessed 31-Aug-2022].