

LLM-based policy generation for intent-based management of applications

Kristina Dzeperoska*

*Department of Electrical and Computer Engineering
University of Toronto
Toronto, Ontario
kristina.dzeperoska@mail.utoronto.ca*

Ali Tizghadam

*Department of Electrical and Computer Engineering
University of Toronto
Toronto, Ontario
ali.tizghadam@utoronto.ca*

Jieyu Lin*

*Department of Electrical and Computer Engineering
University of Toronto
Toronto, Ontario
jieyu.lin@mail.utoronto.ca*

Alberto Leon-Garcia

*Department of Electrical and Computer Engineering
University of Toronto
Toronto, Ontario
alberto.leongarcia@utoronto.ca*

Abstract—Automated management requires decomposing high-level user requests, such as intents, to an abstraction that the system can understand and execute. This is challenging because even a simple intent requires performing a number of ordered steps. And the task of identifying and adapting these steps (as conditions change) requires a decomposition approach that cannot be exactly pre-defined beforehand. To tackle these challenges and support automated intent decomposition and execution, we explore the few-shot capability of Large Language Models (LLMs). We propose a pipeline that progressively decomposes intents by generating the required actions using a policy-based abstraction. This allows us to automate the policy execution by creating a closed control loop for the intent deployment. To do so, we generate and map the policies to APIs and form application management loops that perform the necessary monitoring, analysis, planning and execution. We evaluate our proposal with a use-case to fulfill and assure an application service chain of virtual network functions. Using our approach, we can generalize and generate the necessary steps to realize intents, thereby enabling intent automation for application management.

I. INTRODUCTION

The growing heterogeneous and distributed resources that support the plethora of services and applications can be challenging to manage, in particular considering dynamic environments and strict application requirements such as availability, security, and reliability. Human-based operations are prone to error, and time- and cost-sensitive. Automation is highly desirable to relieve administrators from repetitive and tedious management tasks, thereby keeping OpEx low, and focusing their attention on more complex problems. Recent advances in ML and AI can be leveraged to simplify these management efforts.

An intent defines a set of operational goals (that a system should meet) and outcomes (that a system should deliver), without specifying how to achieve or implement them [1]. Intent-based networking (IBN) or intent-based management

seek to automate network and management tasks by allowing the system to accept and realize user intents. Intents bring two main requirements that are categorized by functionality into intent fulfillment and intent assurance. The former includes: functions such as translation, decomposition and execution; abstractions that capture system hierarchy (from top-level applications to low-level infrastructure devices) and provide a logical view to support the intent-related functions. To ensure that intents are continuously met during their life-cycle, and to prevent "intent drift" [1], intent assurance requires the system to monitor and adapt the intent deployment accordingly. Finally to close the loops for assurance and fulfillment, negative feedback control loops can couple the required steps for the necessary measurements, decisions and control.

Large Language Models (LLMs) provide powerful capabilities for Natural Language Processing (NLP) tasks, such as understanding, generating, and classification of text data. This makes LLMs attractive as a solution in translation and decomposition in intent processing. Moreover, LLM models are trained on massive datasets, which allows a model to learn and understand the contextual relationship in natural language data. This is beneficial for intent-based management because it provides the ability to learn and adapt the steps required to realize intents. In this paper, we leverage the few-shot learning capability of LLMs to generate progressively the steps to realize an intent. This approach can generalize to unseen intents and support dynamic environments, where pre-defined sequences of actions would fail.

We use a policy-based approach to capture and model the relevant abstractions at each level of the management hierarchy. Our goal is to automatically decompose intents into a sequence of policies that when executed will deploy the intent. Our focus is on the capabilities of generic large language models for intent-based application management.

We developed "Emergence", an intent based management system. In order to realize the notion that an intent specifies

*equal contribution

”what the user wishes the system to do, without specifying how”, we implemented an LLM pipeline for progressive decomposition of intents into policies, and then use mapping functions for policy-to-API resolution. Through this pipeline, Emergence determines the ”how” part of the intent. This allows Emergence to then proceed with intent deployment, and to support intent fulfillment and assurance.

We use a language model to understand natural language data and context in order to derive the steps to realize an intent. These steps enable our system to gather the necessary monitoring data, analyze the data based on current environment conditions, and accordingly create and execute a plan to deploy the intent onto the infrastructure. To link these steps and execute the policies, we use feedback control loops (MAPE-K, monitor-analyze-plan-execute and knowledge [2]).

We evaluate Emergence in an intent use-case that involves a virtual network function (VNF) service chain with high availability, consisting of a deep packet inspection tool and a load-balancer. Our use-case demonstrates both intent fulfillment and assurance, by taking appropriate actions when issues are detected to ensure that the intent is continuously met during its life-cycle.

Our contributions in this paper are:

- Propose the use of LLMs with few-shot learning for progressive decomposition of intents into policies to support intent-based application management.
- Evaluate the system on a cloud testbed and demonstrate intent fulfillment and assurance.
- Discuss opportunities of generic LLMs for intent-based management.

II. BACKGROUND AND RELATED WORK

Research in intent-based systems has focused on different requirements, scopes, architectures, and approaches for the intent-related tasks. Efforts by standardization bodies such as Internet Engineering Task Force (IETF), European Telecommunications Standards Institute (ETSI), International Telecommunication Union (ITU), and TM Forum, as well as in academia have considered the use of ML/AI, closed control loops, abstractions (formal intent languages, policy models, etc.). Recent surveys discuss these research efforts, requirements, and challenges [3]–[5].

Natural-language processing (NLP) is required to translate and formalize intents. Recurrent neural networks, such as long short-term memory (LSTM), have been used to translate intents [6]–[10]. These models extract the necessary information from the intent, which is then mapped to an abstraction model or language that the system understands in order to process the intent. For example, LUMI [6] uses bidirectional LSTMs to extract entities and translate intents into an intent language (NILE), and then into configuration commands. LUMI has a chat-based interface (Google Dialogflow) for users to express their intents to the system. Other chat-based proposals include: iNDIRA [11] uses NLP to construct semantic RDF graphs to understand, interact, and create the required network services; EVIAN [12], an extension of iNDIRA, uses the RASA NLP

tool to translate intents. RASA is also used in [13], and Dialogflow is used in [14] and [8].

Large language models (LLMs) such as GPT ([15], [16]), PaLM [17] and LLaMA [18] are transformer-based models with impressive capabilities for NLP tasks. These models are trained using massive amounts of text data, and have billions of parameters. While both LSTMs and LLMs handle context in text, LLMs typically outperform LSTMs in large-scale language tasks due to a self-attention mechanism [19], that provides a more flexible way of handling long-range dependencies in text. Moreover, LLMs are capable of few-shot learning [16], which allows the model to learn to perform a new task by simply training the model with several examples for a given task at hand. ChatGPT by OpenAI uses the Generative Pretrained Transformer (GPT) language model to generate responses to natural language inputs. Although these models are not trained for intent-based management tasks, they can be useful tools for intent processing due to their capabilities for few-shot learning. In [20], we explored this capability to decompose intents into a set of Python APIs. Here, we leverage LLMs for intent to policy decomposition.

Abstraction languages allow the modeling of an intent to a formal model that the system can understand. Languages proposed for intent modeling include, NILE [6], SNIL [8], LAI [21]. We use the formal policy framework from our prior work, [22] that allows us to model policies at different levels of abstractions across the system hierarchy. To promote the adoption of our policy framework, we created a mapping to the Metro Ethernet Forum Policy Driven Orchestration Model (MEF PDO) [23], which is extensive and YANG-defined to support automation.

III. METHODOLOGY

We now describe the Emergence system for management of intent-based applications. In this paper we focus on the LLM pipeline to decompose intents into a policy-based abstraction that can be mapped to APIs for intent execution and deployment. Figure 1 gives an overview of the pipeline and its three stages. Each uses an LLM with few-shot learning. The first stage classifies intents into types known by the system. The second stage decomposes the intent and type into policies. The third stage validates the policies for omissions or errors in the policy format, as well as for the correct ordering of policies. Before discussing the pipeline in detail, we describe our policy-based abstraction.

A. Policy model functional abstraction

We use a policy-based approach to provide the abstractions needed at each level of the system hierarchy. We defined a formal policy framework in [22], to model different types of policies (e.g., utility, goal, action) at different levels of abstractions, and detect and resolve conflicts across the hierarchy.

Our model formally defines Policy \vec{P} as:

$$\vec{P} = (D, E, A, \vec{C}). \quad (1)$$

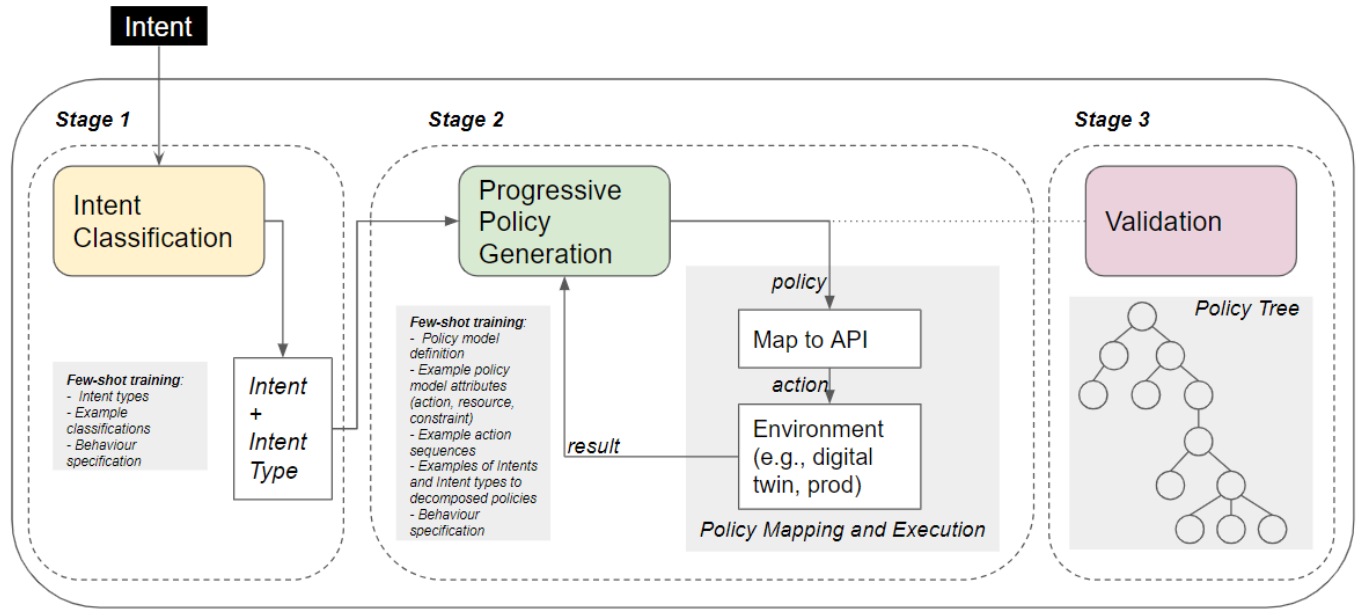


Fig. 1. Pipeline overview: 1) classify intents to known intent types, 2) progressively decompose intents and generate policies: map each policy to an API, execute the API and return the result to the LLM, 3) validate the policies. The resulting policy tree represents the sequence of derived policies from the intent.

where D denotes the policy definer that defines the policy, E is an entity or a group of entities that enforce the policy, and A is an action. \vec{C} is a vector of constraints that apply to action A with regards to resources (\vec{R}), temporal (\vec{T}) and spatial (\vec{S}) constraints, defined as:

$$\vec{C} = (\vec{R}, \vec{T}, \vec{S}). \quad (2)$$

For each policy \vec{P} , we define policy metadata that includes the policy ID, domain, expiration date, priority, and autonomic permission. To support implementation of our policies in order to execute and deploy intents, we have created mapping functions to convert our policies to our APIs and to the MEF PDO model.

B. Intent Decomposition using few-shot learning

Our goal is to decompose intents into policies with actions that correspond to a MAPE-K closed control loop. We use our policy model as an abstraction and train a generic LLM (OpenAI's ChatGPT) using few-shot learning to generate policies for user intents. In the few-shot part, we pass to the LLM an input message that generally contains: definitions (e.g., policy model) and descriptions, examples of intents and their corresponding intent types and policy sequences, and LLM behaviour specification (e.g., to act as a MAPE-K sequential policy generator). We find that LLMs can generalize well and learn to decompose intents, even with just a couple of intent to policy examples.

The generated policies represent the ordered sequence of policies to process and deploy an intent at a given time. We refer to this sequence as a policy tree (an example tree is shown in Figure 1). As an example, for an intent that calls for the creation of a small virtual machine, the following

policy (one of the required policies) is needed to check for availability of 1 small virtual machine in Domain1: $P = (\text{avail}, \text{vm}, \text{zone} = \text{Domain1}, \text{size} = \text{small}, \text{count} = 1)$. To train the LLM we represent the policy as a JSON object as follows: $\{\text{"action"} : \text{"avail"}, \text{"resource"} : \text{"vm"}, \text{"zone"} : \text{"Domain1"}, \text{"size"} : \text{"small"}, \text{"count"} : 1\}$. This allows us to train the LLM to output the policies in the same format (key-value pairs), making it easier to integrate them into code. For example, a dictionary data structure in Python allows us to map the policy action to an API and include the remaining policy attributes. We note that we purposefully omit the Definer and Enforcer attributes from the policies for clarity. The Definer is the user or application that creates or manages the intent, and the Enforcer is the MAPE component that enforces the policy. For example, for the above policy, the Definer is "Administrator" and the Enforcer is "Analyze".

1) *Intent Classification Stage*: The first pipeline stage uses an LLM to classify intents into one or more intent types supported by our system. To train the LLM with few-shot learning, we provided the intent types supported in our system, examples of intents and their corresponding intent types, and LLM behaviour specification (e.g., to act as an intent classifier). An intent may contain one or more intent types. For example, the intent: "Create a small monitored VM in domain 1." is classified to the following types: *create resource*, *schedule health check*. The first type is for the set of policies to create the VM. The second type is for the health check, since a VM cannot be monitored without having this check. Other intent types include: *deploy service*, *start service*, *stop service*, *run service*, *discover resource*, *collect resource*, *publish resource*, *validate resource*, etc. Figure 2 shows the output of the classification results for three different intents.


```

Policy Generator : {"action":"get", "resource":"compute", "zone":"Domain1"}
Calling API for action: get
Result provided to policy generator: True
Policy Generator : {"action":"get", "resource":"flavors", "zone":"Domain1"}
Calling API for action: get
Result provided to policy generator: True
Policy Generator : {"action":"avail", "resource":"vm", "zone":"Domain1", "count":1, "size":["medium", "medium", "small", "small"]}
Calling API for action: avail
Result provided to policy generator: True
Policy Generator : {"action":"reserve", "resource":"vm", "zone":"Domain1", "count":1, "size":["medium", "medium", "small", "small"]}
Calling API for action: reserve
Result provided to policy generator: True
Policy Generator : {"action":"create", "resource":"vm", "zone":"Domain1", "count":1, "name":["Appl_dpi", "Appl_lb", "Appl_web1", "Appl_web2"], "image":["dpi", "lb", "web", "web"], "size":["medium", "medium", "small", "small"]}
Calling API for action: create
Created VM ID and Name: f64b0a1-1960-4997-8f1d-a3f17c576777 Appl_web2
Created VM ID and Name: bd453c26-7348-492f-833d-8435f0e507c0 Appl_web1
Created VM ID and Name: b82f519c-d47b-4c86-a48c-2fc37d3d31a3 Appl_lb
Created VM ID and Name: c8374eb-10ca-4409-ea42-bfa1d0e683ee Appl_dpi
VM IP: ['10.30.11.10', '10.30.11.24', '10.30.11.15', '10.30.11.7']
Result provided to policy generator: True
Policy Generator : {"action":"validate", "resource":["Appl_dpi", "Appl_lb", "Appl_web1", "Appl_web2"]}
Calling API for action: validate
Result provided to policy generator: True
Policy Generator : {"action":"deploy", "resource":["Appl_dpi", "Appl_lb", "Appl_web1", "Appl_web2"], "service":"SFC"}
Calling API for action: deploy
Result provided to policy generator: True
Policy Generator : {"action":"healthcheck", "resource":["Appl_dpi", "Appl_lb", "Appl_web1", "Appl_web2"]}
Calling API for action: healthcheck
Result provided to policy generator: True
Policy Generator : {"action":"schedule", "resource":"healthcheck_policy", "frequency":"hourly"}
Calling API for action: schedule
Result provided to policy generator: True
Policy Generator : {"action":"notify", "resource":"schedule_policy", "recipient":"AppManagement", "status":"issue"}
Calling API for action: notify
Result provided to policy generator: True
Policy Generator : END
Intent deployed

```

Fig. 5. Progressive policy generation and execution for intent fulfillment. Intent: "Deploy a service function chain with high availability in Domain1 consisting of: a medium vm for the dpi service, a medium vm for the load-balancer service, and 2 small vms for the web servers". Intent type: create resource, deploy service, availability.

```

Policy Generator : {"action":"start", "resource":"Appl_dpi", "zone":"Domain1"}
Calling API for action: start
Result provided to policy generator: False
Policy Generator : {"action":"delete", "resource":"Appl_dpi", "zone":"Domain1"}
Calling API for action: delete
Result provided to policy generator: True
Policy Generator : {"action":"get", "resource":"compute", "zone":"Domain1"}
Calling API for action: get
Result provided to policy generator: True
Policy Generator : {"action":"get", "resource":"flavors", "zone":"Domain1"}
Calling API for action: get
Result provided to policy generator: True
Policy Generator : {"action":"avail", "resource":"vm", "zone":"Domain1", "count":1, "size":"medium"}
Calling API for action: avail
Result provided to policy generator: True
Policy Generator : {"action":"avail", "resource":"image", "name":"dpi", "zone":"Domain1"}
Calling API for action: avail
Result provided to policy generator: True
Policy Generator : {"action":"reserve", "resource":"vm", "zone":"Domain1", "count":1, "size":"medium"}
Calling API for action: reserve
Result provided to policy generator: True
Policy Generator : {"action":"create", "resource":"vm", "zone":"Domain1", "count":1, "name":"Appl_dpi", "image":"dpi", "size":"medium"}
Calling API for action: create
Created VM ID and Name: 11d5b5ee-37c1-4941-a35a-993d52038a74 Appl_dpi
VM IP: ['10.30.11.8']
Result provided to policy generator: True
Policy Generator : {"action":"validate", "resource":"Appl_dpi"}
Calling API for action: validate
Result provided to policy generator: True
Policy Generator : {"action":"update", "resource":["Appl_dpi", "Appl_lb", "Appl_web1", "Appl_web2"], "service":"SFC"}
Calling API for action: update
Result provided to policy generator: True
Policy Generator : END
Intent assured

```

Fig. 6. Progressive policy generation and execution for intent assurance (second scenario).

are provided as parameters to the API. Figure 5 shows the result of this stage. In summary, the LLM outputs policies to gather necessary data, check resource availability, make a reservation request, create and validate resources, deploy the service function chain, and creates policies to enable monitoring for high availability (i.e., schedule a health check, and set up notifications). If these actions complete successfully, the intent can be safely deployed (given that the validation stage also completes successfully).

B. Intent Assurance

To demonstrate assurance, we intentionally shut off the DPI VM to trigger the assurance. This was captured by the health check, which sent a notification about the state of the monitored VMs to the Application Manager module (AppManagement). This module detects the intent drift (the state of the DPI VM is "Shutdown", as opposed to "Running"), and passes this message to the LLM to fix the intent. To test the assurance capability of the LLM, we test two scenarios. In the first scenario, we let the LLM successfully perform a start

TABLE I
INTENT EVALUATION ANALYSIS

	Execution Time (s)	Number of policies
Intent Fulfillment	338.9	11
Intent Assurance 1	13.2	2
Intent Assurance 2	85.7	10

action for the DPI VM. In this case, two policies are executed: start the VM, and validate the VM. In the second case, we provide a negative result when trying to perform the start action, and in turn the LLM generates additional assurance-related policies: delete the VM, get necessary data and check resource availability again, make a reservation request, create and validate the VM, and last, update the service function chain. The result of the second scenario is shown in Figure 6.

C. Evaluation Results and Analysis

The results from fulfillment and assurance are shown in Figures 7 and 8 respectively. We report the average times and the number of generated policies to fulfill and assure the intent (for both assurance scenarios) in Table I. The execution results mostly depend on the time our testbed APIs take to complete (e.g., create VM). The time to receive a response using the ChatGPT API is mostly negligible. Although not immediately evident, the policies help us quantify the hidden complexities when working with intents. For example, to fulfill the use-case intent, 11 policies were generated, and these policies are essentially defining the main logic for how to realize the intent. Considering that the policies get mapped to API calls, this means our policies abstract even more lines of code executed through the API calls.

One important aspect of an intent-based system is the ability to generalize well, meaning that the system can handle modified and unseen intents. From our evaluations, we find that the LLM can handle this requirement well. For example, the intent used for the few-shot example was simpler and had different requirements compared to the use-case intent. However, the LLM was able to learn from the provided context (few-shot learning), such that it can generalize and decompose intents with unseen requirements. These results indicate that generic LLMs are very promising for intent-based applications and management.

ID	Name	Status	Task State	Power State	Networks
9e5b74eb-10e3-4409-a1c2-bfe50e4d38aa	Appl_dpi	ACTIVE	-	Running	sc-kristina-net=10.30.11.7
5831915c-da7b-4888-485c-2fe241db81a2	Appl_lb	ACTIVE	-	Running	sc-kristina-net=10.30.11.18
5d513c26-7340-452f-833d-8435fce907c2	Appl_web1	ACTIVE	-	Running	sc-kristina-net=10.30.11.24
4e9a50e1-1900-4197-921a-82d17d574677	Appl_web2	ACTIVE	-	Running	sc-kristina-net=10.30.11.10

Fig. 7. Result of the intent fulfillment.

ID	Name	Status	Task State	Power State	Networks
1145b5ee-9701-4841-a95a-599d52038a74	Appl_dpi	ACTIVE	-	Running	sc-kristina-net=10.30.11.8
5831915c-da7b-4888-485c-2fe241db81a2	Appl_lb	ACTIVE	-	Running	sc-kristina-net=10.30.11.18

Fig. 8. Result of the intent assurance.

V. DISCUSSION AND OPPORTUNITIES

We now discuss our insights and additional LLM opportunities for intent-based applications, lessons learned, limitations, and future work.

A. LLMs for additional conversions

In this paper, we demonstrated that LLMs can be used to decompose intents into a high-level abstraction, such as a policy. Per our policy framework, the policies specify the actions, resources, and constraints to be applied. However, to execute these policies, we either need to map them to a language that supports implementation, such as YANG (e.g., to the MEF PDO policy model), or to a corresponding API. In this work, we use separate mapping functions in our implementation to map policies into APIs. However, an LLM can also be trained to be able to convert policies into other formal languages, including programmable code, e.g., Python APIs. The benefit of deriving the policies first, is that it allows us to obtain the intent workflow, i.e., the set of steps to fulfill or assure intents in a MAPE loop.

In our prior work (Emergence), we use a set of pre-defined Finite State Machines (FSM) to execute policies. However, it would be more beneficial to be able to dynamically generate these FSMs. In this manner, the system would be able to execute new, unseen intents, provided that the required functions (e.g., APIs) are available. One approach to creating dynamic, on-demand FSMs, would be to train a generic LLM to create Boolean-Logic Decision Trees from the policies, and then convert these trees into Finite State Machines.

B. Lessons Learned

An important consideration for LLMs with few-shot learning include over-training and under-training. It is worth noting that we do not modify the weights of the model, but instead provide descriptions, definitions, and usage examples as input to the model, so that the underlying Transformer architecture can learn to decompose intents into policies. The goal is to be able to train the model to generalize well, in order to consistently generate policies that are correct and work as intended. To help generalize better, before a new intent is received, we clear the conversation history, and we send the training prompt to the model. If the training prompt contains too many details, it can confuse the LLM, or it can cause the LLM to output exactly as per the prompt (even if incorrect). If there is too little information, the LLM can begin to deviate and create new attributes or entities (e.g., actions, constraints, resources) that are irrelevant to the system.

In this paper, we provide simple results back to the LLM upon the execution of a policy, for example, a Boolean value (True, False). If more information is provided to the model, such as some details from the execution, then the LLM could modify the policy constraints accordingly in an attempt to meet the intent. For example, if a policy asks for a large size VM, or a specific version of an image, and these are not available, then, if the LLM receives information about other available sizes and images as part of the result, the LLM will select

the best option, e.g., a medium VM, and an image version as close as possible to the original intent version requested. This shows that we can further improve the reasoning logic for the LLM, for example by embedding specific algorithms as part of the few-shot training. In this way, the LLM can play an even bigger role in management.

C. Limitations and Future Work

Although we demonstrated that a generic large language model can be very efficient in decomposing intents using few-shot learning, there are a number of challenges to consider:

- Validation is important to ensure that the policies capture the desired behaviour. This requires validating the policies in terms of the format, proper sequence, correctness and omissions of attributes. For this, we can train an LLM to look for these types of issues. Second, we need to check for logic and algorithmic issues. For this, we can leverage FSMs to ensure correct state transitions, and that the final state matches the intent's desired state. Finally, it is desirable to use a combination of simulators and emulators to create a real digital twin environment.
- The transformer model comes with context length limitation, i.e., the number of tokens that the model can attend to. This limits how far the transformer can look back in the conversation history. As such, this also limits the size of the training prompt. However, recent models are increasing the context limit, e.g., GPT3.5 supports 4096 tokens, and GPT4 variants offer 8k or 32k.
- The Boolean results provided to the LLM can inadvertently cause the LLM to relax some important constraints, or even start generating policies that deviate from the user's intent (i.e., beyond the scope) as the LLM tries to "desperately" meet the intent. Therefore, it is important to carefully craft the intent, for example by using specific pointer words such as *shall*, *must*, *should* to "limit" the LLM's efforts. For example, consider an intent to use a path between two nodes with some specific requirements. If such a path is not available currently, the LLM might try to use actions such as engineering the traffic, or even the network, if the previous attempt fails. This is a major overreaction to a simple intent.

VI. CONCLUSIONS

In this paper, we focused on leveraging a generic LLM in our Emergence system for intent-based management and applications. Specifically, we proposed using LLMs with few-shot learning to enable progressive policy generation, driven by policy execution results. We considered both intent fulfillment and assurance, and our results indicate that LLMs are very promising in attaining our goal to enable automatic intent decomposition for application management.

REFERENCES

- [1] A. Clemm, L. Ciavaglia, L. Z. Granville, and J. Tantsura, "Intent-Based Networking - Concepts and Definitions," RFC 9315, Oct. 2022. [Online]. Available: <https://www.rfc-editor.org/info/rfc9315>
- [2] J. Kephart and D. Chess, "The vision of autonomic computing," *Computer*, vol. 36, no. 1, pp. 41–50, 2003.
- [3] A. Leivadeas and M. Falkner, "A survey on intent-based networking," *IEEE Communications Surveys Tutorials*, vol. 25, no. 1, pp. 625–655, 2023.
- [4] E. Zeydan and Y. Turk, "Recent advances in intent-based networking: A survey," in *2020 IEEE 91st Vehicular Technology Conference (VTC2020-Spring)*, 2020, pp. 1–5.
- [5] L. Pang, C. Yang, D. Chen, Y. Song, and M. Guizani, "A survey on intent-driven networks," *IEEE Access*, vol. 8, pp. 22 862–22 873, 2020.
- [6] A. S. Jacobs, R. J. Pfitscher, R. H. Ribeiro, R. A. Ferreira, L. Z. Granville, W. Willinger, and S. G. Rao, "Hey, lumi! using natural language for {intent-based} network management," in *2021 USENIX Annual Technical Conference (USENIX ATC 21)*, 2021, pp. 625–639.
- [7] Y. Ouyang, C. Yang, Y. Song, X. Mi, and M. Guizani, "A brief survey and implementation on refinement for intent-driven networking," *IEEE Network*, vol. 35, no. 6, pp. 75–83, 2021.
- [8] M.-T.-A. Nguyen, S. B. Souihi, H.-A. Tran, and S. Souihi, "When nlp meets sdn : an application to global internet exchange network," in *ICC 2022 - IEEE International Conference on Communications*, 2022, pp. 2972–2977.
- [9] C. Yang, X. Mi, Y. Ouyang, R. Dong, J. Guo, and M. Guizani, "Smart intent-driven network management," *IEEE Communications Magazine*, vol. 61, no. 1, pp. 106–112, 2023.
- [10] N. Vedula, N. Lipka, P. Maneriker, and S. Parthasarathy, "Open intent extraction from natural language interactions," in *Proceedings of The Web Conference 2020*, 2020, pp. 2009–2020.
- [11] M. Kiran, E. Pouyoul, A. Mercian, B. Tierney, C. Guok, and I. Monga, "Enabling intent to configure scientific networks for high performance demands," *Future Generation Computer Systems*, vol. 79, pp. 205–214, 2018.
- [12] H. Mahtout, M. Kiran, A. Mercian, and B. Mohammed, "Using machine learning for intent-based provisioning in high-speed science networks," in *Proceedings of the 3rd International Workshop on Systems and Network Telemetry and Analytics*, 2020, pp. 27–30.
- [13] C. H. Cesila, R. P. Pinto, K. S. Mayer, A. F. Escallón-Portilla, D. A. A. Mello, D. S. Arantes, and C. E. Rothenberg, "Chat-ibn-rasa: Building an intent translator for packet-optical networks based on rasa," in *2023 IEEE 9th International Conference on Network Softwarization (NetSoft)*, 2023, pp. 534–538.
- [14] M. Bezahaf, E. Davies, C. Rotsos, and N. Race, "To all intents and purposes: Towards flexible intent expression," in *2021 IEEE 7th International Conference on Network Softwarization (NetSoft)*, 2021, pp. 31–37.
- [15] OpenAI, "Gpt-4 technical report," 2023.
- [16] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell *et al.*, "Language models are few-shot learners," *Advances in neural information processing systems*, vol. 33, pp. 1877–1901, 2020.
- [17] A. Chowdhery, S. Narang, J. Devlin, M. Bosma, G. Mishra, A. Roberts, P. Barham, H. W. Chung, C. Sutton, S. Gehrmann *et al.*, "Palm: Scaling language modeling with pathways," *arXiv preprint arXiv:2204.02311*, 2022.
- [18] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar *et al.*, "Llama: Open and efficient foundation language models," *arXiv preprint arXiv:2302.13971*, 2023.
- [19] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [20] J. Lin, K. Dzevaroska, A. Tizghadam, and A. Leon-Garcia, "Appleseed: Intent-based multi-domain infrastructure management via few-shot learning," in *2023 IEEE 9th International Conference on Network Softwarization (NetSoft)*. IEEE, 2023, pp. 539–544.
- [21] B. Tian, X. Zhang, E. Zhai, H. H. Liu, Q. Ye, C. Wang, X. Wu, Z. Ji, Y. Sang, M. Zhang *et al.*, "Safely and automatically updating in-network acl configurations with intent language," in *Proceedings of the ACM Special Interest Group on Data Communication*, 2019, pp. 214–226.
- [22] K. Dzevaroska, N. Beigi-Mohammadi, A. Tizghadam, and A. Leon-Garcia, "Towards a self-driving management system for the automated realization of intents," *IEEE Access*, vol. 9, pp. 159 882–159 907, 2021.
- [23] "MEF Standard (95), Policy Driven Orchestration (PDO)," Metro Ethernet Forum, July 2021. [Online]. Available: <https://www.mef.net/wp-content/uploads/MEF-95.pdf>