

AGENP: An ASGrammar-based GENerative Policy Framework

Seraphin Calo*, Irene Manotas*, Dinesh Verma*, Elisa Bertino†, Mark Law‡, Alessandra Russo‡

*IBM Research

Yorktown Heights, NY

scalco@us.ibm.com, irene.manotas@ibm.com, dverma@us.ibm.com

†Purdue University

W. Lafayette, IN, USA

bertino@purdue.edu

‡Imperial College London

South Kensington Campus, London

mark.law09@imperial.ac.uk, a.russo@imperial.ac.uk

Abstract—Generative policies have been proposed as a mechanism to autonomously adapt to changes in a system to achieve the systems’s goals with minimum to none human intervention. The combination of generative policies and learning mechanisms can help a coalition system to be more effective when working in a distributed, continuously transforming environment with a diverse set of members, resources, and tasks. Learning mechanisms based on logic e.g., Inductive Logic Programming (ILP), have several properties that make them suitable and attractive for the creation and adaptation of generative policies. Some of the ILP features include representation of the data and problems in a declarative way, the ability to learn a general model with a small number of examples, and with an existing background knowledge. Recently, Answer Set Programming (ASP) have been recognized as a powerful language for knowledge representation and reasoning for which ILP has been successfully applied to.

This paper proposes AGENP, a ASGrammar-based GENerative Policy Framework for Autonomous Managed Systems (AMS) that aims to support the creation and evolution of generative policies by leveraging ILP. We describe the framework components, the required inputs, data structures, and mechanisms to support the refinement and instantiation of policies, identification of policy violations, and monitoring of policies and policy adaptation according to changes in the AMS and its context. We also present the main workflow for the global and local refinement of policies and their adaptation based on ASP and ILP concepts and tools, and present example scenarios where the AGENP framework can be useful for the self-generation of policies.

Keywords-policy-based management systems; generative policies; inductive logic programming

I. INTRODUCTION

With the recent advances in Artificial Intelligence (AI), Autonomous Systems (AS) composed of intelligent, self-managed devices are on the rise [4, 13]. Traditional policy management tools need to evolve in order to be able to manage a myriad of complex situations where autonomous devices are operating, in isolation or collaboratively, to ensure that devices and systems take secure and appropriate decisions. Enabling policies to be autonomously adapted

and created according to dynamic contexts, where resources availability and conditions change over time, could help autonomous managed systems to take timely decisions, while maintaining consistency and compliance with the system objectives.

Policy-based Management Systems (PBMSs) usually work with predefined policies that are provided by admin users, or experts on the system’s operations, to ensure the proper system functionality. Policies definitions are usually not modified unless a human analyze them and pinpoint corrections or improvements, and new policies are only provided to the system by end users when new business objectives are identified. This manual approach for policy creation and modification is tedious and error prone [11]. Hence, approaches for guided and automatic policy generation have been proposed [11, 12, 16]. However, these approaches are either focused on static policies (i.e., predefined policies that do not change over time) or they rely on large amounts of data to be able to identify new policies, or required policies’ transformations. To the best of our knowledge, there has not been approaches that try to learn a policy model to self-generate policies that capture common and optimal policy decisions of a managed system.

The Generative Policy-based Model (GPM), recently proposed in [15], proposes the concept of *generative policies* i.e., policies that are self-generated by the managed system (e.g., devices) to enable systems to determine their own behavior, allowing flexible policy definitions that make managed systems more autonomous. A generative policy model could be learned from an initial specification of policies of the managed system, and from the analysis of the system’s operations, resource usage, and previous policy decisions. One learning approach that can be used for learning generative policy models is by leveraging Inductive Logic Programming (ILP) techniques. ILP is a learning approach that have been successful at learning representative models (e.g., programs) from examples and background knowledge [5]. ILP is an appealing symbolic learning ap-

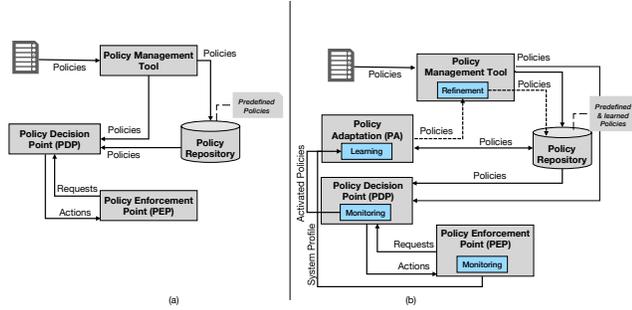


Figure 1. Policy-based Management Systems (PMS): (a) Standard PMS, (b) Evolved PMS.

proach because, contrary to other popular supervised learning approaches, it is a white-box or knowledge oriented method that provides easy to understand models [7], and does not require large amounts of labeled data to generate a good model of the underlying system’s knowledge. This paper presents the ASGrammar-based GENERative Policy Framework (AGENP), a PBMS framework based on GPM and that is designed to learn a generative policy model via ILP with the goal of supporting autonomous management systems in self-generating and selecting optimal policies based on system’s operation and history of policy decisions.

II. EVOLUTION OF POLICY-BASED MANAGEMENT SYSTEMS

A. Standard Policy-based Management Systems

Fig. 1 (a), shows a standard PBMS, such as the one defined by IETF/DMTF¹, where policy-based management tools create, edit, and control policies stored in a persistent policy repository. A Policy Decision Point (PDP) interprets policies in the policy repository and sends decisions to a Policy Enforcement Point (PEP) using an appropriate format. Then, the PEP executes an action, or set of actions, according to the decisions obtained from the PDP. Policies used in these systems are usually provided by administrator users or user experts on the system’s functionality. In these PBMS, changes to policies do not occur unless an admin user analyzes the current policies and identifies those that need to be corrected or improved, and new policies are provided by a domain expert to the managed system when new business goals are identified.

B. Evolved Policy-based Management Systems

Changes in the internal and external conditions of the managed system can create scenarios where initial predefined policies cannot be used because such situations represent events and conditions not foreseen initially in the policies provided to the PBMS. Learning from the history of actions taken by the managed system, analyzing

policies’ decisions and policies’ structure (i.e., event, conditions, actions), and improving policy definitions, would make the managed system to be prepared to make decisions for situations not initially anticipated, thus enabling its autonomous operation. Contrary to the standard policy-based management system, the evolved policy-based management system includes support for *policy refinement* and *policy learning* to allow learning from past experiences, and from the continuous evolution of policies. Figure 1 (b) shows an overview of the evolved policy-based management system that includes the policy refinement and policy learning processes described in general below and in detail in Section IV and Section V.

Policy Refinement: The refinement process allows to find the best policy representation for policies used by the Autonomous Managed System (AMS). The policy refinement process continuously consolidate the set of policies, and their representations, according to the policy constraints initially established for the system, and based on decisions made by the managed system during its operation.

Policy Learning: Learning from previous policy decisions allows the managed system to identify new policies that can be used to handle unforeseen situations while complying with the system objectives and constraints, and without waiting for policies to be discovered and provided by an end user. The policy learning process learns a policy model that considers the refined policies, the decisions made by the system in the past, and the system context (e.g., system status and environment). The policy model is used to instantiate new policies that are predicted to be useful for the immediate and future system’s operations. Section IV describes in detail the structure of the GPM that is designed to support operations of autonomous managed systems, and Section V describes the AGENP framework, which is an instantiation of the Generative Policy-based Model (GPM) using ILP as the learning approach.

III. LEARNING POLICY MODELS

We define the process of identifying the general policy characteristics i.e., objects, conditions, and actions, and their associated values, from a set of policies, and for a given policy domain type (e.g., access-control policy or network policy), as learning a policy model. A policy model can be represented as a grammar, or policy template, that can be used to instantiate policies for a given domain of an AMS. We propose inductive learning as a policy learning approach to capture a policy model for a set of policies, that can be used to refine, instantiate, and adapt policies according to the operations of a AMS.

A. Inductive Logic Programming (ILP)

In ILP the main task is to find an hypothesis, that together with given background knowledge, explains a set of observations (i.e., positive examples). ILP is appealing for

¹<https://tools.ietf.org/html/rfc2753>

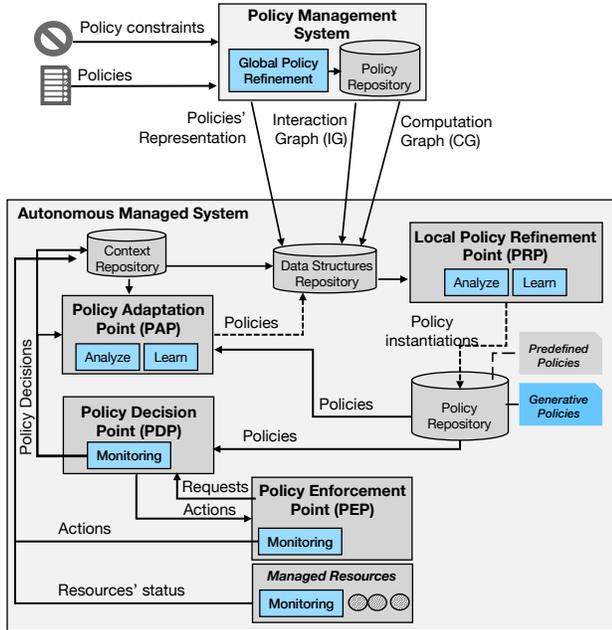


Figure 2. Overview of the Generative Policy Model Architecture

learning tasks mainly because its ability to learn a general model with a small number of examples and with an existing background knowledge [8], and also because the learned model (or hypotheses) can be expressed in plan english making them easier for a human user to understand [10].

Answer Set Program (ASP) is a declarative programming paradigm based on logic programming, under answer set semantics, that have been used extensively in reasoning applications, and to support the development of inference engines. Recently, an ILP approach for ASP was proposed [8], which allows to learn inductive solutions as ASP. We propose using ASP, and Answer Set Grammar (ASG) as a low-level representation of policies, because ASP is a highly declarative paradigm for knowledge representation and reasoning, and it allows to specify policy specifications such as the ones required by the generative policy definition [2].

Next section presents a review of the GPM, and Section V describes the design of the ASGrammar-based GENERative Policy Framework (AGENP) that uses ASP and ILP techniques to learn a generative policy model to instantiate policies according to the GPM model.

IV. THE GENERATIVE POLICY MODEL (GPM)

This section presents a more detailed description of the GPM [15], including inputs and outputs of the model, and the components located in the management tool and in the managed system. Figure 2 shows an overview of the components for the policy-based management of au-

tonomous managed systems in the generative policy model architecture.

A. Inputs and Outputs

The GPM takes as inputs *policy definitions* and *policy constraints*. Policy definitions are provided by admin users to reflect the correct actions to take for different conditions of the managed system's operation. Policy constraints allow the AMS to establish bounds within which policies can be defined. Policy boundaries can be seen as regions of the search space of policy instantiations that are should never be explored. Policy constraints serve as a mechanism to bound the size of the policy search space. The outputs consist of the policy representations for the policies that have been found to be optimal given the conditions of the AMS operation, and the policy violations that have been found during the analysis and learning of policy decisions.

B. GPM Components

Beyond the traditional components of a PBMS, which include a PDP, a PEP, and a policy repository, the GPM follows the general architecture of the evolved PBMS described in Section II-B. The GPM includes two new components, a Policy Adaptation Point (PAP) at the AMS, and a Policy Refinement Point (PRP) at the management tool and at the AMS. Also, two additional repositories are defined for the AMS, the context information repository and the data structures repository, which are described below along with the new components.

1) *Global Policy Refinement*: Policy refinement is done at the global level, i.e., in the management system, and at the local level, inside the autonomous managed system, in a Policy Refinement Point (PRP) as shown in Fig. 2. Global policy refinement allows the policy management system to consider policies that are within the prescribed bounds, and to obtain the right representation for the policies (i.e., grammar or template). In the case that the management system is collaborative i.e., multiple managed systems agree to collaborate to perform tasks together, then global refinement is in charge of defining the Interaction Graph (IG) and the Computation Graph (CG) that specify the roles and tasks that can be assigned to the different resources of the managed systems involved in the collaborative scenario.

2) *Data Structures Repository*: The data structures repository stores the policy representations (grammar), policy constraints, and interaction graph, along with their versions i.e., initial version specified by the admin user, and the subsequent versions generated by the AMS during its operation.

3) *Context Information Repository*: The context information repository holds information about the status of the management system. The monitoring of decisions in the PDP and the status of resources in the PEP is stored in this repository for future reference. The information keep in the context repository helps in the refinement and adaptation

can be used to instantiate policies. Policy instantiations are saved in the policy repository so that the AMS can use them for its operations. The generated 'revised' attribute grammar is shared back to the ASP subcomponent. Once there has not been changes between consecutive versions of the grammar being generated the 'revised' attribute grammar is stored in the data structures repository.

Policy Adaptation Point (PAP): The policy adaptation process uses the latest attribute grammar from the policy data repository. History of policy decisions and context information, such as resources' status and availability, are used as examples and background knowledge by the analyzer, where positive examples can be policies with high quality (i.e., policies that are consistent, complete, minimal, relevant, and correct [6]) or. The attribute grammar is converted into an ASG in the ASP component and then it is passed to the ASP Solver to find the optimal set of policy solutions along with policy violations. Policy solutions are then passed to a learner for ASP, such as the ILASP2 learner [9], that is able to learn a preference ordering of the generative policy model. Both the preference ordering and the revised grammar are stored in the data structures repository.

VI. GENERATIVE POLICY SCENARIOS

Various scenarios involving Autonomous Managed System could take advantage of the generative policies concept and a framework instantiation like AGENP. This section highlights some scenarios that we think would benefit from instantiations of the GPM in general.

A. Policies for Federated Learning Environments

Federated systems, i.e., systems where multiple organizations work together to share resources and collaborate to achieve common tasks, could benefit from using generative policies to manage its operations as shown in [1]. In the same way, a Federated Learning (FL) system, where the main objective is to learn a common Machine Learning (ML) model, requires the definition and management of policies for distributed learning and data/model fusion. These systems can use generative policies to manage and ensure the right operations of the FL system are in place. FL-related policies need to consider different characteristics of the FL system and its environment [14], such as the quality of the data and communications, the data arrangements among the different nodes, the type of ML algorithm, and the performance of different participating nodes, to decide how to distribute the data and learn a federated ML model. Given the wide diversity of learning environments and algorithms that can exist in FL systems, a framework like AGENP that helps to self-generate policies could provide great advantages for the proper functionality of the FL system by identifying over time which policies are optimal, and how policies can evolve to handle operations between different AMS working together.

B. Policies for Autonomous Vehicles Operations

Autonomous transportation and systems for autonomous vehicles are becoming more common than before [3]. Policies that help to improve the user experience, i.e., focusing on personalization and context awareness by providing the ability for the vehicle to adapt to the behavior and desires of the users in a safe and secure way, are an important aspect of autonomous vehicles. Generative policies in this context can help autonomous vehicles to identify which personalization policies, based on the vehicle context information and user preferences, should be used during the vehicle operation based on the user, the vehicle status, and its environment. Also, generative policies can be used to pinpoint policies that define how more complex autonomous systems interact, for instance for situations where multiple autonomous vehicles are operating at the same time, interacting with each other.

VII. RELATED WORK

Autonomic management is an area where policies are used to automatically manage a system under predefined conditions. In this area tools like Madison [11], a library for automatic policy generation for SELinux policies, have been developed to provide modules that guide the policy definition process, and modules that extract information from reference policy interface files to allow the creation of policy templates from system's accesses. However, to the best of our knowledge, none of the current autonomic management approaches for consider the evolution of policies according to the commons decisions made by the managed system. The closest work related to ours is the one done by Quiroz et al. that presents a framework for autonomic policy adaptation based on online analysis and characterization of system operation based on decentralized clustering techniques [12]. Our approach is different to theirs in that we use symbolic learning to capture common and optimal policies based on the system operation without relying on large amounts of data but on a small set of examples and background knowledge about the managed system policies, its components, and operation. Jabal et al. presents a review of policy analysis approaches and tools. Policy examples for access control (e.g., role-based access control policies), and network (e.g., firewall policies, SDN policies) are presented, as well as the policy analysis process, and the frameworks and tools available [6]. As we have presented in this paper, different policy types and analysis tools can be leveraged by other instantiations of the GPM, and in the AGENP framework to analyze and select optimal polices' decisions for different AMS and scenarios.

REFERENCES

- [1] E. Bertino, S. Calo, M. Toma, D. Verma, C. Williams, and B. Rivera. A cognitive policy framework for next-generation distributed federated systems: Concepts and

- research directions. In *IEEE 37th International Conference on Distributed Computing Systems (ICDCS)*, pages 1876–1886, 2017.
- [2] Seraphin Calo, Dinesh Verma, Supriyo Chakraborty, Elisa Bertino, Emil Lupu, and Gregory Cirincione. Self-generation of access control policies. In *Proceedings of the 23rd ACM on Symposium on Access Control Models and Technologies, SACMAT*, pages 39–47, 2018.
- [3] Riccardo Coppola and Maurizio Morisio. Connected car: Technologies, issues, future trends. *ACM Comput. Surv.*, 49(3):46:1–46:36, October 2016.
- [4] C. L. Fok, M. Hanna, S. Gee, T. C. Au, P. Stone, C. Julien, and S. Vishwanath. A platform for evaluating autonomous intersection management policies. In *3rd International Conference on Cyber-Physical Systems*, pages 87–96, 2012.
- [5] Sumit Gulwani, José Hernández-Orallo, Emanuel Kitzelmann, Stephen H. Muggleton, Ute Schmid, and Benjamin Zorn. Inductive programming meets the real world. *Commun. ACM*, 58(11):90–99, 2015.
- [6] A. A. Jabal, M. Davari, E. Bertino, C. Makaya, S. Calo, D. Verma, A. Russo, and C. Williams. Methods and tools for policy analysis. 2018.
- [7] Dimitar Kazakov and Daniel Kudenko. Multi-agents systems and applications. chapter Machine Learning and Inductive Logic Programming for Multi-agent Systems, pages 246–270. Springer-Verlag New York, Inc., 2001.
- [8] Mark Law, Alessandra Russo, and Krysia Broda. Inductive learning of answer set programs. In Eduardo Fermé and João Leite, editors, *Logics in Artificial Intelligence*, pages 311–325. Springer International Publishing, 2014.
- [9] Mark Law, Alessandra Russo, and Krysia Broda. Learning weak constraints in answer set programming. *Theory and Practice of Logic Programming*, 15(4-5): 511–525, 2015.
- [10] Mark Law, Alessandra Russo, and Krysia Broda. The complexity and generality of learning answer set programs. *Artificial Intelligence*, 259:110–146, 2018.
- [11] K. Macmillan. Madison : A new approach to policy generation? In *SELinux Symposium*, 2007.
- [12] Andres Quiroz, Manish Parashar, Nathan Gnanasambandam, and Naveen Sharma. Autonomic policy adaptation using decentralized online clustering. In *Proceedings of the 7th International Conference on Autonomic Computing, ICAC*, pages 151–160. ACM, 2010.
- [13] Dinesh Verma, Graham Bent, and Ian Taylor. Learning neural network policies with guided policy search under unknown dynamics. In *Proceedings of the 9th International Conference on Advanced Cognitive Technologies and Applications, COGNITIVE*, 2017.
- [14] Dinesh Verma, Seraphin Calo, and Greg Cirincione. Distributed AI and security issues in federated environments. In *Proceedings of the Workshop Program of the 19th International Conference on Distributed Computing and Networking, Workshops ICDCN*. ACM, 2018.
- [15] Dinesh Verma, Seraphin Calo, Supriyo Chakraborty, Elisa Bertino, C. Williams, J. Tucker, and B. Rivera. Generative policy model for autonomic management, 2017.
- [16] Le Yu, Tao Zhang, Xiapu Luo, and Lei Xue. Autoppg: Towards automatic generation of privacy policy for android applications. In *Proceedings of the 5th Annual ACM CCS Workshop on Security and Privacy in Smartphones and Mobile Devices, SPSM*, pages 39–50. ACM, 2015.