Learning Intuitive Grammars for Movement Primitive Sequences

Rudolf Lioutikov[1], Guilherme Maeda[2], Filipe Veiga[1], Kristian Kersting[3] and Jan Peters[1,4]

***Abstract*— Movement Primitives are a well studied and widely applied concept in modern robotics. Composing primitives out of an existing library, however, has shown to be a challenging problem. We propose probabilistic context-free grammars to sequence a series of primitives to generate complex robot policies from a given library of primitives. The rule-based nature of formal grammars allows an intuitive encoding of hierarchically and recursively structured tasks. This hierarchical concept strongly connects with the way robot policies can be learned, organized, and re-used. The induction of context-free grammars is often formulated as a maximum a-posteriori optimization. Given the common formulation of the prior, the resulting grammars are often not easily comprehensible by non-experts. We introduce a novel, easy to tune prior aimed at producing intuitive grammars for movement primitive sequencing.**

Movement primitives (MPs) are a well established concept in robotics. MPs are used to represent atomic, elementary movements and are, therefore, appropriate for tasks consisting of a single stroke-based or rhythmic movement [1]. However, for more complex tasks a single MP is often not sufficient and a sequences of MPs is required. A declared goal of robotics is the deployment of robots into scenarios where direct or indirect interactions with non-expert users are required. Therefore, intuitive sequencing mechanisms for non-experts are necessary.

This work proposes probabilistic context-free grammars (PCFGs) for the sequencing of MPs and introduces a novel prior to induce PCFGs from observed sequences of primitives. PCFGs have been intensively studied in both natural language processing and compiler construction but have also been applied in a variety of fields, e.g., molecular biology, bioinformatics, computer vision and robotics [2], [3]. PCFGs allow the implicit embedding of hierarchies within the rules of the grammar associating every produced sequence with at least one corresponding parse tree. Such a parse tree represents the derivation of the produced sequence in an intuitive way. Figure 1 shows a learned grammar for placing a stone in a game of tic-tac-toe, including the parse tree for a produced primitive sequence. However, the understandability of the grammar itself depends on the size of both the grammar, i.e., the number of possible productions, as well as the length of each possible production. The induction of concise but expressive grammars is considered non-trivial.

A common approach to grammar induction is to formulate the problem as a maximum-a-posteriori estimation [4], where the scoring is defined as the posterior of the grammar $\mathcal{G}$ given the observations $\mathcal{D}$

[1]Intelligent Autonomous Systems, TU Darmstadt, 64289 Darmstadt, Germany, {lioutikov, veiga, peters}@ias.tu-darmstadt.de
[2]ATR Computational Neuroscience Labs, 2-2-2 Hikaridai Seika-sho, Soraku-gun, Kyoto 619-0288, Japan, g.maeda@atr.jp
[3]CS Dept. and Centre for Cognitive Science, TU Darmstadt, 64289 Darmstadt, Germany, kersting@cs.tu-darmstadt.de
[4]Max Planck Institute for Intelligent Systems, 72070 Tübingen, Germany

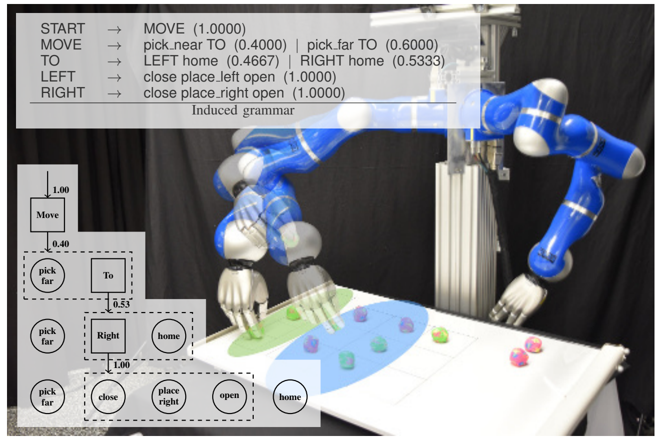

Fig. 1: The robot executes a turn in the tic-tac-toe game, represented as a sequence of movement primitives. The sequence was generated by a probabilistic context-free grammar learned from previously labeled observations.

$$\begin{aligned}\mathcal{G}^* &= \arg\max_{\mathcal{G}} p\left(\mathcal{G}|\mathcal{D}\right), \\ &= \arg\max_{\mathcal{G}} p\left(\mathcal{D}|\,\mathcal{G}\right) p\left(\mathcal{G}\right).\end{aligned} \tag{1}$$

The likelihood is commonly defined as

$$p\left(\mathcal{D}|\,\mathcal{G}\right) = \prod_{\boldsymbol{d}\in\mathcal{D}} \frac{1}{|\mathrm{parse}\left(\boldsymbol{d},\mathcal{G}\right)|} \sum_{\tau\in\mathrm{parse}(\boldsymbol{d},\mathcal{G})} \prod_{\rho\in\tau} \rho, \tag{2}$$

where $\boldsymbol{d}$ denotes an observed sequence of primitives and $\mathrm{parse}\left(\boldsymbol{d},\mathcal{G}\right)$ represents the set of all possible parse trees for the sequence $\boldsymbol{d}$ given the grammar $\mathcal{G}$. Each parse tree τ consists of the rules that were involved in the production of $\boldsymbol{d}$, with ρ being the probability of a single rule.

Several grammar priors $p\left(\mathcal{G}\right)$ have been suggested, separating the prior into a parameter and a structure term $p\left(\mathcal{G}\right) = p\left(\boldsymbol{\rho}_{\mathcal{G}}\middle|\,\mathcal{G}_{\mathcal{R}}\right) p\left(\mathcal{G}_{\mathcal{R}}\right)$. The parameter term $p\left(\boldsymbol{\rho}_{\mathcal{G}}\middle|\,\mathcal{G}_{\mathcal{R}}\right)$ is then represented as a Dirichlet distribution over the grammar parameters $\boldsymbol{\rho}_{\mathcal{G}}$ and the structure prior $p\left(\mathcal{G}_{\mathcal{R}}\right)$ is usually defined as an exponential distribution over the Minimal Description Length (MDL) of the grammar.

However, an uninformative Dirichlet distribution will very quickly dominate the likelihood and favor grammars with a high number of rules per nonterminal. Given that the MDL serves as the energy, the structure prior will favor small grammars. In combination this quickly results in grammars with many short rules, not necessarily intuitively understood by non-experts.

We introduce a novel grammar prior aimed at favoring grammar structures that are more comprehensible by non-experts. We define the prior jointly over the parameters and

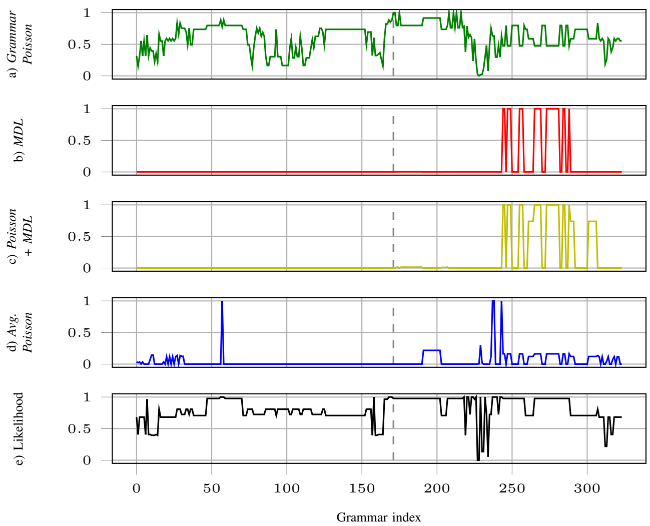

Fig. 2: The posteriors and the likelihood for the tic-tac-toe turn grammar. The dashed line indicates the index of the highest posterior (171), given the presented Poisson prior.

the structure $p(\mathcal{G}) = p\left(\boldsymbol{\rho}_{\mathcal{G}}, \mathcal{G}_{\mathcal{R}}\right)$,

$$p\left(\boldsymbol{\rho}_{\mathcal{G}}, \mathcal{G}_{\mathcal{R}}\right) = \frac{p\left(|\mathcal{R}||\eta_{\mathcal{R}}\right)}{|\mathcal{R}|} \sum_{(\mathrm{R}_A, \boldsymbol{\rho}_A) \in \mathcal{R}} p\left(|\mathrm{R}_A||\eta_{\mathrm{R}}\right) p\left(\mathrm{R}_A|\boldsymbol{\rho}_A, \eta_r\right). \quad (3)$$

where the probabilities over the number of rules $p\left(|\mathcal{R}||\eta_{\mathcal{R}}\right)$ and the size of each rule $p\left(|\mathrm{R}||\eta_{\mathrm{R}}\right)$ are modeled as Poisson distributions with means $\eta_{\mathcal{R}}$ and η_{R}. The probability of each rule is modeled as a weighted average

$$p\left(\mathrm{R}_A|\boldsymbol{\rho}_A, \eta_r\right) = \sum_{r \in \mathrm{R}_A, \rho \in \boldsymbol{\rho}_A} \rho \; p\left(|r||\eta_r\right), \quad (4)$$

over the probabilities of the corresponding productions. The weighting is given by the grammar parameters $\rho \in \boldsymbol{\rho}_A$ and the probability of each production corresponds to the Poisson distribution over its length $p\left(|r||\eta_r\right)$, given a desired production length η_r. Since all components are defined as discrete probabilities, the prior is always $p(\mathcal{G}) \leq 1$, yielding a stronger influence of the likelihood and, hence, the observations in the posterior than common grammar priors. Furthermore, the prior $p(\mathcal{G})$ will now prefer grammars with η_{R} productions per nonterminal with an average length of η_r symbols per production. The hyper-parameters $\eta_{\mathrm{R}}, \eta_r$ can be set to achieve a desired simplicity of the grammar. By weighting each production $r \in \mathrm{R}_A$ with the corresponding grammar parameter $\rho \in \boldsymbol{\rho}_A$ the prior gives more significance to productions that are more likely to occur.

We evaluated the novel prior on a tic-tac-toe pick-and-place task and compared the resulting posterior, *Grammar Poisson*, with the ones from three common structure prior choices, *MDL* [5], *Poisson + MDL* [6], *Avg. Poisson* [3], based on the MDL and/or a Poisson distribution over the length of each rule. A major difference of the *Grammar Poisson* prior to the other discussed priors is that we do not model the distribution over the grammar parameters as a Dirichlet distribution but rather use them as a weighting for the average production length. Each produced sequence corresponds to one turn of the game, i.e. picking a stone, closing the hand, placing the stone on the field , opening the hand and returning to the home position.

The grammar learning was initialized with 16 observations of four unique sequences, each consisting of five terminals. We initialized the approach with $\eta_{\mathcal{R}} = 5$, $\eta_{\mathrm{R}} = 2$ and $\eta_r = 3$. The corresponding normalized posteriors are shown in Figure 2a and the grammar with the highest posterior, grammar index 171, as well as a possible sequence including the corresponding parse tree, are shown in Figure 1. Figure 2b-d shows the normalized posteriors corresponding to the three common priors. The x-axis corresponds to the different grammars traversed during the optimization. The spiky behavior of the posteriors (b-d) is due to the uninformative Dirichlet prior for the grammar parameters and the exponential distribution over the MDL. Both of these factors can change significantly with a small change in the grammar. It is noticeable that the likelihood of the grammar $p\left(\mathcal{G}^i|\mathcal{D}\right)$ does not play significantly into the posteriors of (b-d), whereas our posterior (a) shows a much stronger dependency on the likelihood. This behavior, is explained by the fact that the likelihood as introduced in Equation (2) is a probability mass function, but the three priors (*MDL*, *Poisson + MDL*, *Avg. Poisson*) are products of probability density functions. In contrast, our prior (*Grammar Poisson*) is defined as a probability mass function, averaging over multiple Poisson distributions. This definition prohibits the prior from completely dominating the likelihood. As a consequence, the proposed prior (*Grammar Poisson*) results in a posterior (a) that takes the given observations stronger into account than the posteriors in (b-d).

In this work, we introduced probabilistic context-free grammars as a mechanism to sequence movement primitives. The paper presents an approach to induce such grammar from flat sequences of movement primitive samples, i.e., no hierarchy in the observations. The new introduced grammar prior is defined over several coupled Poisson distributions, and eliminates complications that arise from both Dirichlet parameter priors and MDL based structure priors. In our method, the hyper-parameters of the prior have an clear semantic interpretation, namely the number of productions for each nonterminal and the average length of each production. The learned grammars are simple and intuitive as evaluated on a real robot platform.

[1] Paraschos, Daniel, Peters, and Neumann, “Using probabilistic movement primitives in robotics,” *Autonomous Robots (AURO)*, accepted.

[2] Dantam and Stilman, “The motion grammar: Analysis of a linguistic method for robot control,” *IEEE Trans. Robotics*, vol. 29, pp. 704–718, 2013.

[3] Lee, Su, Kim, and Demiris, “A syntactic approach to robot imitation learning using probabilistic activity grammars,” *Robotics and Autonomous Systems*, vol. 61, pp. 1323–1334, 2013.

[4] Stolcke, “Bayesian learning of probabilistic language models,” Ph.D. dissertation, Berkeley, CA, USA, 1994, uMI Order No. GAX95-29515.

[5] Talton, Yang, Kumar, Lim, Goodman, and Mech, “Learning design patterns with bayesian grammar induction,” in *25th Annual Symposium on User Interface Software and Technology, October 7-10, 2012, Cambridge, MA, USA,*, Miller, Benko, and Latulipe, Eds. ACM, 2012, pp. 63–74.

[6] Kitani, Sato, and Sugimoto, “Recovering the basic structure of human activities from noisy video-based symbol strings,” *IJPRAI*, vol. 22, pp. 1621–1646, 2008.