

Geometric Multi-Model Fitting by Deep Reinforcement Learning

Zongliang Zhang, Hongbin Zeng, Jonathan Li,* Yiping Chen, Chenhui Yang, Cheng Wang

Fujian Key Laboratory of Sensing and Computing for Smart Cities, School of Information Science and Engineering
Xiamen University, Xiamen, Fujian 361005, China (*Corresponding author)

{zhangzongliang, hbzeng}@stu.xmu.edu.cn, {junli, chenyping, chyang, cwang}@xmu.edu.cn

Abstract

This paper deals with the geometric multi-model fitting from noisy, unstructured point set data (e.g., laser scanned point clouds). We formulate multi-model fitting problem as a sequential decision making process. We then use a deep reinforcement learning algorithm to learn the optimal decisions towards the best fitting result. In this paper, we have compared our method against the state-of-the-art on simulated data. The results demonstrated that our approach significantly reduced the number of fitting iterations.

Introduction

Geometric model fitting aims to reconstruct underlying models (e.g., lines, circles, characters, and buildings) from given data (e.g., images or laser scanning point clouds). With the reconstructed rich model information (e.g., shape, scale, rotation, and location), the data can be comprehensively understood. With such merit, model fitting has constantly attracted research interests for a long time. However, the model fitting problem is far from being solved, at least in terms of computational speed, because of increasing complexity of encountered data and thus models. A common case of complex data is that data conceive multiple models. For example, a CAPTCHA image usually contains multiple characters (George et al. 2017). A multi-model fitting technique is needed to handle such data.

A recent trend for addressing model fitting problem is to formulate it as an optimization problem (Lake, Salakhutdinov, and Tenenbaum 2015), such that it can be conveniently tackled by an existing optimization algorithm. Our previous method (Zhang et al. 2019) uses the cuckoo search (CS) algorithm (Yang and Deb 2010) to solve the optimization problem in model fitting, and notably achieves the new state-of-the-art in the challenging few-shot character recognition tasks (George et al. 2017). CS can approach the optimum with high precision. However, it usually takes many iterations to converge to the optimum, especially when the fitting involves a large number of variables, which is the case of multi-model fitting. The number of variables involved in n -model fitting is as large as n times that in single-model fitting. In other words, it is time-consuming to use CS to perform multi-model fitting.

Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

In this paper, we propose a reinforcement learning approach for optimization in multi-model fitting. Our insight is as follows. The selection of variable values for a model can be seen as a decision. The fitting of multiple models is a process that consists of a sequence of decisions. Such decision making process can be efficiently optimized by reinforcement learning.

The work similar to ours can be found in (Teboul et al. 2013), which uses a traditional reinforcement learning method to control binary split shape grammar for parsing facade images. Their method works under the assumption that the split grammar has discrete variables. However, the variables involved in model fitting usually are continuous, which are challenging for a traditional reinforcement learning method to handle (Lillicrap et al. 2015). In contrast, our work is based on recently developed deep reinforcement learning (DRL), which has made remarkable progress for a number of challenging tasks including continuous control (Lillicrap et al. 2015).

Method

A geometric model M is a k -dimensional point set, i.e., $M \subset \mathbb{R}^k$. In this paper, $k = 3$. Give a data point set $D \subset \mathbb{R}^k$, the goal of model fitting is to find a model M that is most similar to D . For multi-model fitting, M is the union set of multiple models, i.e., $M = \bigcup_{i=1}^n M_g^{\theta_i}$, where n is the number of models, and M_g^{θ} is the model defined by a given parametric rule g which is parameterized by variable θ . Formally, multi-model fitting can be formulated as the following maximization problem:

$$\max_{(\theta_1, \theta_2, \dots, \theta_n)} f(\theta_1, \theta_2, \dots, \theta_n) = s(\bigcup_{i=1}^n M_g^{\theta_i}, D), \quad (1)$$

where $s(\cdot, \cdot)$ is the geometric similarity estimator defined in (Zhang et al. 2019). The pseudo-code of our method is shown in Algorithm 1, where the DRL actor q is based on (Lillicrap et al. 2015), and \mathcal{N}_i is exploration noise (Lillicrap et al. 2015). Our method follows a hypothesis and verify paradigm to solve the maximization problem Eq. (1). In each iteration, for each model, a hypothesis value θ_i is proposed according to the actor and exploration noise. Then the hypothesis is verified through computing the reward in order to update the actor.

Algorithm 1 The proposed method

input: a data point set D , a parametric rule g with variable θ , the number of models n , the DRL actor q , the verify function f

output: $(\theta_1^*, \theta_2^*, \dots, \theta_n^*)$ that maximizes Eq. (1)

Randomly initialize $(\theta_1^*, \theta_2^*, \dots, \theta_n^*)$

for iteration $j = 1$ to j_{max} **do**

for $i = 1$ to n **do**

 Select variable $\theta_i = q(i) + \mathcal{N}_i$

 Observe reward:

$$r_i = f(\theta_1, \theta_2, \dots, \theta_i) - f(\theta_1, \theta_2, \dots, \theta_{i-1})$$

 Update q according to r_i

end for

if $f(q(1), q(2), \dots, q(n)) > f(\theta_1^*, \theta_2^*, \dots, \theta_n^*)$ **then**

$$(\theta_1^*, \theta_2^*, \dots, \theta_n^*) = (q(1), q(2), \dots, q(n))$$

end if

end for

We now present the computation cost of the proposed DRL based n -model fitting method. In each iteration, the computational time cost of a hypothesis and verify algorithm is composed of two parts: the hypothesis part t_{DRL}^H and the verify part t_{DRL}^V . Therefore, the total computational cost is $t_{DRL} = j_{DRL}(t_{DRL}^H + t_{DRL}^V)$, where j_{DRL} is the total number of iterations. For n -model fitting, it is needed to calculate the verify function f for n times in each iteration. Let t_f be the cost to calculate f one time, then $t_{DRL}^V = nt_f$, and $t_{DRL} = j_{DRL}(t_{DRL}^H + nt_f)$. In contrast, the CS based method (Zhang et al. 2019) only needs to calculate f one time in one iteration. Consequently, the total computational cost of the CS based method is $t_{CS} = j_{CS}(t_{CS}^H + t_f)$. It can be concluded that, when $t_f \gg t_{DRL}^H$ and $j_{CS} > nj_{DRL}$, DRL is more efficient than CS. Note that t_f is determined by the data and the model sizes (Zhang et al. 2019). Therefore, $t_f \gg t_{DRL}^H$ holds for many applications in which data and model sizes are large. For example, a laser scanning point cloud is usually large in size as containing millions of points.

Experiments

In this abstract, we preliminarily evaluate our method by fitting line segments to the data D_2 shown in Fig. 1b. Specifically, the parametric rule g input to our method is a vertical line segment rule with only one variable $\theta \in \mathbb{R}$ that determines the horizontal location of the line segment. We also fix the number of models $n = 4$.

As shown in Figs. 1c and 1d, DRL fits the data well after only 100 iterations, whereas CS cannot well fit the data even after 1000 iterations. The evolutions of similarity during fitting are shown in Fig. 2, where each line represents the mean of similarity values and the patches around each line represents the standard deviations. The mean values and standard deviations are computed from 5 times of fitting. The results clearly indicates that DRL is tens of times more efficient than CS in terms of the numbers of fitting iterations.




Figure 1: Data point sets and fitted models. (a) Clean data D_1 . (b) Corrupted data D_2 , which is generated by adding some outliers to D_1 . (c) The model fitted by CS after 1000 iterations. (d) The model fitted by DRL after 100 iterations.

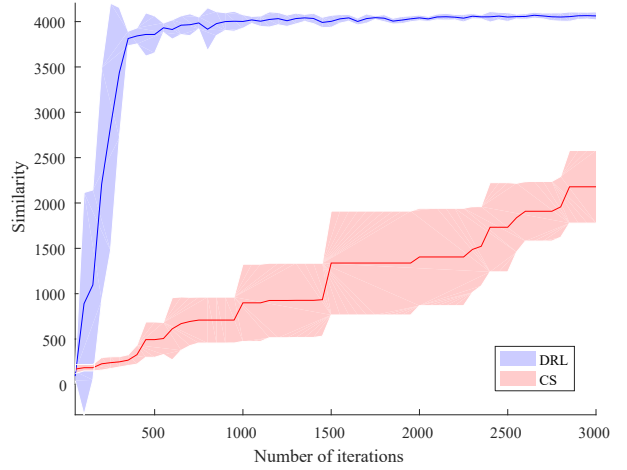


Figure 2: Fitting similarities at different iterations of DRL and CS.

References

- George, D.; Lehrach, W.; Kansky, K.; Lazaro-Gredilla, M.; Laan, C.; Marthi, B.; Lou, X.; Meng, Z.; Liu, Y.; Wang, H.; et al. 2017. A Generative Vision Model That Trains with High Data Efficiency and Breaks Text-Based CAPTCHAs. *Science* 358(6368):eaag2612.
- Lake, B. M.; Salakhutdinov, R.; and Tenenbaum, J. B. 2015. Human-Level Concept Learning Through Probabilistic Program Induction. *Science* 350(6266):1332–1338.
- Lillicrap, T. P.; Hunt, J. J.; Pritzel, A.; Heess, N.; Erez, T.; Tassa, Y.; Silver, D.; and Wierstra, D. 2015. Continuous Control with Deep Reinforcement Learning. *arXiv preprint arXiv:1509.02971*.
- Teboul, O.; Kokkinos, I.; Simon, L.; Koutsourakis, P.; and Paragios, N. 2013. Parsing Facades with Shape Grammars and Reinforcement Learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35(7):1744–1756.
- Yang, X. S., and Deb, S. 2010. Engineering Optimisation by Cuckoo Search. *International Journal of Mathematical Modelling and Numerical Optimisation* 1(4):330–343.
- Zhang, Z.; Li, J.; Guo, Y.; Li, X.; Lin, Y.; Xiao, G.; and Wang, C. 2019. Robust Procedural Model Fitting with a New Geometric Similarity Estimator. *Pattern Recognition* 85:120 – 131.