International Journal of Approximate Reasoning ••• (••••) •••–•••

Contents lists available at ScienceDirect

International Journal of Approximate Reasoning

www.elsevier.com/locate/ijar

Quick and energy-efficient Bayesian computing of binocular disparity using stochastic digital signals

Alexandre Coninx *, Pierre Bessière, Jacques Droulez

ISIR CNRS/UPMC, 4 place Jussieu, 75005 Paris, France

ARTICLE INFO

Article history:
Received 15 May 2016
Received in revised form 31 October 2016
Accepted 1 November 2016
Available online xxxx

Keywords:
Bayesian inference
Stochastic computing
Sensory processing
Energy efficiency
Hardware implementation
Binocular disparity

ABSTRACT

Reconstruction of the tridimensional geometry of a visual scene using the binocular disparity information is an important issue in computer vision and mobile robotics, which can be formulated as a Bayesian inference problem. However, computation of the full disparity distribution with an advanced Bayesian model is usually an intractable problem, and proves computationally challenging even with a simple model. In this paper, we show how probabilistic hardware using distributed memory and alternate representation of data as stochastic bitstreams can solve that problem with high performance and energy efficiency. We put forward a way to express discrete probability distributions using stochastic data representations and perform Bayesian fusion using those representations, and show how that approach can be applied to diparity computation. We evaluate the system using a simulated stochastic implementation and discuss possible hardware implementations of such architectures and their potential for sensorimotor processing and robotics.

1. Introduction

Using two cameras in a stereoscopic setup to reconstruct the tridimensional geometry of a visual scene, in a way similar to that performed by human stereopsis, is an important issue in computer vision, with major applications to autonomous robotics (and more specifically autonomous driving [1]). That issue has been an active research topic since at least 40 years, and a wide range of methods and algorithms have been proposed [2,3] and evaluated on standardized benchmarks [4,5].

Several works have shown that the binocular disparity computation can efficiently be formulated as a Bayesian inference problem [6,7]. The disparity value for each pixel is then expressed as a discrete probability distribution, which can be computed through a probabilistic model using likelihood values specified from the image data. However, computing the full disparity distribution on whole images proves challenging and computationally demanding. That's why most works on binocular disparity using Bayesian models instead reduce the output to a single disparity value per pixel (often using the maximum a-posteriori likelihood estimator). That approach simplifies the computation and allows to reformulate it as an energy minimization problem that can be solved efficiently by classic optimization techniques such as dynamic programming [see 6, for an example].

However, it means that although the computation is based on a probabilistic formalism, it yields deterministic disparity values and not disparity distributions, despite the latter representation being richer and offering many benefits, especially

* Corresponding author.
E-mail address: alexandre.coninx@isir.upmc.fr (A. Coninx).

http://dx.doi.org/10.1016/j.ijar.2016.11.004
0888-613X/© 2016 Elsevier Inc. All rights reserved.

for robotics and sensorimotor systems. Full disparity distributions can accurately represent cases where stereopsis is not sufficient to completely determinate the world geometry, such as ambiguous pixels with multiple matches, or pixels with no matches (e.g. due to occlusions). Such probabilistic representations can also directly be used by Bayesian mapping and navigation methods such as the Bayesian occupation filter [8], and more generally by probabilistic and Bayesian robotics techniques [9–11]. Bayesian inference also provides a powerful framework to express assumptions and prior knowledge about the structure of the world (for example the location of the ground or other known objects) as prior probability distributions.

Stochastic computing is a field dedicated to designing and using computing devices that are intentionally stochastic to perform probabilistic reasoning, using non-Von Neumann architectures, distributed memory and specific data representations. More specifically, the BAMBI project is a research effort to develop stochastic machines implementing Bayesian inference (Bayesian machines) [12]. In this paper, we show how those Bayesian machines can be used to efficiently compute full binocular disparity distribution, paving the way towards fully stochastic autonomous robots and other sensorimotor systems.

In the remainder of this article, we will first give an overview of the related work in section 2, both about stochastic computing and fast binocular disparity computation. We will then describe our Bayesian binocular disparity computation model in section 3. Section 4 will be dedicated to the description of the stochastic computer implementing that model, focusing first on the general principles of computation using stochastic bitstream and second to their application to the Bayesian disparity computation. The evaluation of that system and its results will be presented in section 5 and further discussed in section 6. We will then conclude in section 7 by summing up the implications of that work for the design of Bayesian robotic systems using stochastic components and discussing the future prospects of that topic.

2. Previous work

2.1. Hardware stochastic computing

The general idea of stochastic computations with temporal coding can be traced back to the seminal works of Von Neumann [13] and Gaines [14] who highlighted the interest of such data representations, but their approaches were not widely pursued due to the rapid development of more efficient deterministic computers. The topic has recently received a renewed attention due to the development of probabilistic and Bayesian models in computer science and engineering – and more specifically for sensorimotor and cognitive systems – and the limitations of classic computers to implement those models.

The idea of developing hardware dedicated to Bayesian reasoning has recently been pursued by several teams [15–17], exploring different computational paradigms to perform probabilistic inference. To address the problem of approximate inference Mansinghka [16] uses sampling methods for approximate inference and in a similar way, Jonas designed Markov Chain Monte Carlo based algorithms to provide a representation of probability distributions as sets of samplers [17]. To compute exact inference, a number of different frameworks and toolsets have been put forward. Vigoda [15] designed architectures based on probabilities represented by analog signals, and used the message passing algorithm to compute exact inference. More recently, a research project conducted at the Nanoscale Computing Fabrics Laboratory has led to the design of an unconventional hardware architecture based on electro-magnetic computations to perform inference on Bayesian Network models [18]. Ferreira et al. [19] also showed that exact inference can be efficiently computed using GPU hardware for some high-dimensional problems. Finally, the approach taken by Thakur et al. [20] is quite similar to ours: they use stochastic bitstreams and target special inference problems. They have proposed two frameworks, BEAST (Bayesian Estimation And Stochastic Tracker) and BIND (Bayesian INference in DAG), to perform inference using stochastic electronics on two types of Bayesian models, Hidden Markov Models and Direct Acyclic Graphs (DAG) respectively.

In the framework of the BAMBI project, another stochastic architecture has been proposed to perform naive Bayesian fusion using Muller C-Elements [21], which achieves exact inference with normalization for binary random variables, but create harmful correlations in the stochastic signals and can't be easily extended to non-binary discrete distributions. Other recent work conducted within the BAMBI project has proposed using digital signals with temporal coding to perform Bayesian inference, and a proof-of-concept to solve a simple sensorimotor problem has been put forward [22]. In this paper, we apply the same principles to a more computationally challenging Bayesian model to highlight their benefits.

2.2. Disparity computation

As it provides a way to estimate the depth information using data from standard digital cameras, the binocular disparity problem has received a wide attention since the beginnings of computer vision. Existing approaches have been summarized in reviews [2,3], which show that most methods follow the same general structure which can be divided in three steps:

1. Computing a *matching cost*, which is a positive value associated to each possible pair of matching pixels.[1] The matching cost is a dissimilarity measure: the least likely the pixels are to match, the higher it is. The cost is computed locally,

[1] Most algorithms use rectified image pairs, which allows to only consider pixels on corresponding rows for matching, and limit the disparity to a maximum value D_{max} corresponding to a minimum distance. D_{max} depends on image resolution, camera focal length and visual environment; typical values are 50 to 100 pixels.