Preface

Probabilistic graphical models, as these words suggest, are theoretical constructs whose definition and analysis closely interrelate concepts from graph theory and concepts from probability theory. Graphs are used as effective ways of representing definite relations between the several basic units or elemental variables involved in such a model. Typically, these are relations of a topological character, concerning the reciprocal positions of the units in a spatial context, and may give rise to undirected graphs, directed graphs, or graphs of a mixed type. Probabilities intervene in the specification of local soft constraints, which express dependences within small subsets of the whole set of variables in a model. A network of such local constraints, suitably assembled, determines a probability distribution on the whole set of variables, from which special probability functions can be derived that will serve as guides for probabilistic inferences. These are inferences from evidence concerning some variables for hypotheses concerning some other variables participating in the same model.

Since at least the 1980s, a reader can note a steadily increasing interest in this kind of model by a number of researchers in vision science, especially in the branch known as computer vision. Some reasons that explain a priori such interest are easily seen. First, the visual scenes are spatial realities; the topological relations coded in a graph and concerning definite units in a scene may play crucial roles in its organisation. Second, the formation of a visual scene, from the optical stimuli to the perceptual appearance, may be viewed as an inference process, subject to conditions of uncertainty and ambiguity; the probabilistic form of a model is generally suitable for dealing with these aspects. Third, research on probabilistic graphical models has constructed efficient and reliable methods for finding optimal solutions to inference problems, in terms of finding stationary points of suitable objective functions; those methods may also be of use in vision science, for finding solutions to problems concerning properties or aspects of visual scenes. Fourth, a key idea underlying the probabilistic graphical models is that of representing large range dynamics, concerning extensive sets of variables, in terms of a number of interrelated local processes, each acting on few variables and with a comparatively simple form; similar ideas are considered by vision scientists, in discussing how several elemental factors that have local fields of action may contribute to the generation of a global well-balanced perceptual result. As an addition to this last comment, I remark that possible connections between the research on vision and the theory of field dynamics were considered well before the emergence of probabilistic graphical models in their current form. They were considered, for example, within the frame of the Gestalt theory of visual perception (Köhler, 1920, 1940).

My research experiences are in mathematical psychology and psychophysical experiments on visual perception. I approached the literature of probabilistic graphical models in computer vision with the expectation of finding cues that could be instructive for formal modelling in the psychology of human perception. On perusing the literature, I confirmed the correspondences mentioned above, but I also met with a situation more complex than I had expected. The inventory of inference tasks that,

in computer vision, are addressed in terms of that kind of model is various and includes tasks that, in psychology, would be categorised as problems of cognition, not of perception. Although classified into standard categories — in particular, Markov random fields and Bayesian networks — the models available look much different from one another, not only in the inference tasks addressed but also in the nature of the variables involved, the local constraints applied to the variables, and the overall graphical structures. All this provides a practical demonstration of the high flexibility of the paradigm in question. Further, setting aside the mathematical and computational aspects, there are semantical and ontological questions that demand attention, concerning the meaning and empirical status of the variables, the origin and justification of the local constraints applied to the variables, and the kinds of relations that may be represented by the graphical structure of a model. In addition, there are differences in the verbal and symbolic terms used for equivalent concepts, which contribute to the complexity of the subject.

Given this experience, I judged it worthwhile to prepare a review monograph on the subject, with the following aims: to describe the essentials of the probabilistic graphical approach to vision modelling, using a unitary and explicitly defined system of terms that are suitable both for Markov random fields and for Bayesian networks; to contribute to the elucidation of the semantical and ontological questions that are implicit in that approach, based on lessons that come from the psychological study of vision; to give precise ideas of the methods used for the construction of the models in question and for their application as guides for inferences concerning vision; and to describe representative examples of models belonging to each of the two main classes now mentioned. Given these characteristics, my present work is not an exhaustive review of the available literature. Rather, it is intended to highlight the theoretical and methodological skeleton of the scientific area in question, accompanied by a few salient examples. I hope that it may serve as an introduction to this special branch of vision science and possibly constitute a useful link between the realisations of computer vision and the expectations of perceptual psychology, as concerns the formal modelling of vision abilities.

The monograph is composed of eight chapters, which can be divided into four groups. In Chapters 1–3 I present the basic concepts by distinguishing between the input, output, and principled parts of the information in a vision inference task and characterising the variables and constraints in which those parts are expressed within a model. In Chapters 4 and 5 the essentials of Markov random fields and of Bayesian networks are separately described; in particular, the formation of a global probability distribution from a network of local soft constraints is discussed. In Chapters 6 and 7 general concepts are presented concerning the construction of probabilistic graphical models and the methods used for performing probabilistic inferences guided by them. Lastly, in Chapter 8 four salient examples from the literature are described, two of the Markov random field type and two of the Bayesian network type. Each chapter includes some figures, which illustrate optical realities, the graphical structures of models, or logical relationships between concepts involved in the models and the organisation of the corresponding processes.

The last remark here is about the symbolic notation. A conspicuous character-

istic of the models I am concerned with is the number of terms they involve: variables of different kinds and their ranges of possible values, sites on which the variables are located, relations between the sites, parameters and their spaces, constraints acting on variables and parameters, real-valued functions expressing such constraints, and so on. In expressing by symbols a formal construct implicit in a model, if all the terms involved were made explicit, quite a complex symbolic expression would be obtained, which could be an obstacle to the intelligibility of the arguments developed. For this reason, I adopt the rule of making explicit only the terms that are really important in the particular context in which a construct occurs, keeping all other terms tacit. This implies that the same construct, in distinct contexts, could be presented in slightly different notations, depending on the specific arguments developed in those contexts.

Padova, September 2026 L. B.