

Meta-Interpretive Learning from noisy images

Stephen Muggleton¹  · Wang-Zhou Dai² · Claude Sammut³ ·
Alireza Tamaddoni-Nezhad⁴ · Jing Wen⁵ · Zhi-Hua Zhou²

Received: 31 March 2017 / Accepted: 5 April 2018 / Published online: 7 May 2018
© The Author(s) 2018

Abstract Statistical machine learning is widely used in image classification. However, most techniques (1) require many images to achieve high accuracy and (2) do not provide support for reasoning below the level of classification, and so are unable to support secondary reasoning, such as the existence and position of light sources and other objects outside the image. This paper describes an Inductive Logic Programming approach called Logical Vision which overcomes some of these limitations. LV uses Meta-Interpretive Learning (MIL) combined with low-level extraction of high-contrast points sampled from the image to learn recursive logic programs describing the image. In published work LV was demonstrated capable of high-accuracy prediction of classes such as *regular polygon* from small numbers of images where Support Vector Machines and Convolutional Neural Networks gave

Editors: James Cussens and Alessandra Russo.

✉ Stephen Muggleton
s.muggleton@imperial.ac.uk

Wang-Zhou Dai
daiwz@nju.edu.cn

Claude Sammut
claudio@cse.unsw.edu.au

Alireza Tamaddoni-Nezhad
a.tamaddoni-nezhad@surrey.ac.uk

Jing Wen
wjing@sxu.edu.cn

Zhi-Hua Zhou
zhouzh@nju.edu.cn

¹ Department of Computing, Imperial College London, London, UK

² LAMDA Group, Nanjing University, Nanjing, China

³ School of Computer Science and Engineering, University of New South Wales, Sydney, Australia

⁴ Department of Computer Science, University of Surrey, Guildford, UK

⁵ School of Computer and Information Technology, Shanxi University, Taiyuan, China

near random predictions in some cases. LV has so far only been applied to noise-free, artificially generated images. This paper extends LV by (a) addressing *classification noise* using a new noise-tolerant version of the MIL system Metagol, (b) addressing *attribute noise* using primitive-level statistical estimators to identify sub-objects in real images, (c) using a wider class of background models representing classical 2D shapes such as circles and ellipses, (d) providing richer learnable background knowledge in the form of a simple but generic recursive theory of light reflection. In our experiments we consider noisy images in both natural science settings and in a RoboCup competition setting. The natural science settings involve identification of the position of the light source in telescopic and microscopic images, while the RoboCup setting involves identification of the position of the ball. Our results indicate that with real images the new noise-robust version of LV using a single example (i.e. one-shot LV) converges to an accuracy at least comparable to a thirty-shot statistical machine learner on both prediction of hidden light sources in the scientific settings and in the RoboCup setting. Moreover, we demonstrate that a general background recursive theory of light can itself be invented using LV and used to identify ambiguities in the convexity/concavity of objects such as craters in the scientific setting and partial obscuration of the ball in the RoboCup setting.

Keywords Inductive logic programming · Meta-interpretive learning · High-level vision · Robocup

1 Introduction

Galileo's *Sidereus Nuncius* (Galilei 2004) describes the first ever telescopic observations of the moon. Using sketches of shadow patterns Galileo conjectured the existence of mountains containing hollow areas (i.e. craters) on a celestial body previously thought perfectly spherical. His reasoned description, derived from a handful of observations, relies on a knowledge of (i) classical geometry, (ii) straight line movement of light and (iii) the Sun as an out-of-view light source. This paper investigates the use of Inductive Logic Programming (ILP) (Muggleton et al. 2011) to derive logical hypotheses, related to those of Galileo, from a small set of real-world images. Figure 1 illustrates part of the generic background knowledge used by ILP for interpreting object convexity in Experiment1 (Sect. 5.1).

Figure 1a shows an image of the crescent moon in the night sky, in which convexity of the overall surface implies the position of the Sun as a hidden light source beyond the lower right corner of the image. Figure 1b shows an illusion in which assuming a light source in the lower right leads to perception of *convex* circles on the leading diagonal. Conversely, a light source in the upper left implies their being *concave*. Figure 1c shows how interpretation of a *convex* feature, such as a mountain, comes from illumination of the *right* side of a convex object. Figure 1d shows that perception of a *concave* feature, such as a crater, comes from illumination of the *left* side. Figure 1e shows how Prolog background knowledge encodes a simple recursive definition of the reflected path of a photon.

This paper explores the phenomenon of knowledge-based perception using an extension of Logical Vision (LV) (Dai et al. 2015). In the previous work LV was shown to accurately learn a variety of polygon classes from artificial images with low sample requirements compared to statistical learners. LV generates logical hypotheses concerning images using an ILP technique called Meta-Interpretive Learning (MIL) (Muggleton et al. 2015; Cropper and Muggleton 2016).

Contributions of this paper The main contributions of this paper are:

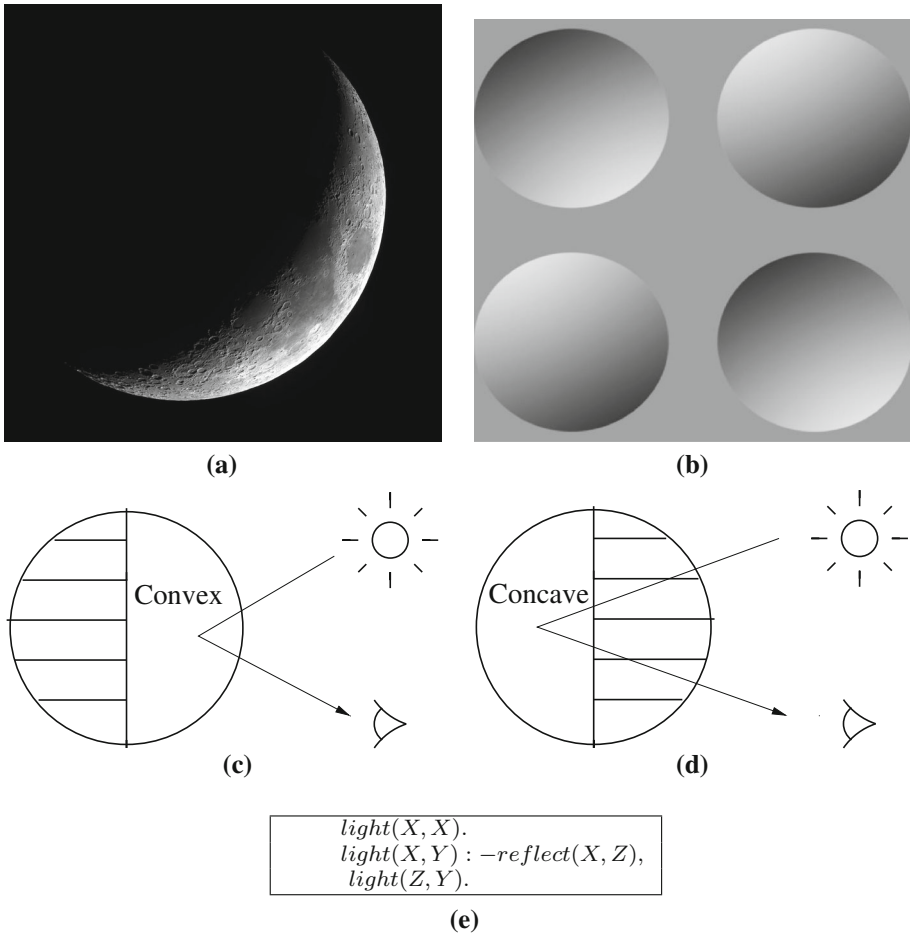


Fig. 1 Interpretation of light source direction: **a** waxing crescent moon (Credit: UC Berkeley), **b** concave/convex illusion, **c** concave and **d** convex photon reflection models, **e** prolog recursive model of photon reflection

1. We describe a generalisation of LV (Dai et al. 2015), which is tolerant to both *classification noise* and *attribute noise*.
2. We show that even in the presence of noise in images [absent in artificial images in Dai et al. (2015)] effective learning can be achieved from as few as one image.
3. We demonstrate that in all cases studied the combination of a logic-based learner with a statistical estimator requires far fewer images (sometimes one) to achieve accuracies requiring large numbers of images using statistical machine learning on its own.
4. We demonstrate that LV can use, as well as invent, generic background knowledge about reflection of photons in providing explanations of visual features.
5. We demonstrate that LV has potential in real application domains such as RoboCup.

RoboCup domain In Experiment 2 (Sect. 5.2) we investigate LV in the context of robotics. Figure 2 shows images from the RoboCup Soccer Standard Platform League.¹ This is a

¹ <http://www.tzi.de/spl/bin/view/Website/WebHome>.




Fig. 2 Robot's view of: **a** another robot and ball clearly separated, **b** the ball partially occluded by a robot, **c** the ball within the bounds of a robot

competition with five Aldebaran Nao robots on each team. They are placed on a $9\text{ m} \times 6\text{ m}$ field, and operate autonomously to play soccer. The robots use cameras to detect the ball, field lines, goals and other robots. In Fig. 2a, the ball can be seen distinctly, whereas in Fig. 2b, c the ball is partially occluded. The problem with recognising the ball is that it consists of several patches of black and white, but there are many other objects on the field that also contain white regions. However, background knowledge concerning the geometry of a sphere projected on a 2D plane guarantees a ball has a circular appearance. If three edge points can be found our approach can fit them to a circle and if that circle has the proportions of black and white pixels, the system concludes it is a ball.

The paper is organised as follows. Section 2 describes related work. The theoretical framework for LV is provided in Sect. 3. Section 4 describes the implementation of LV, including the recursive background knowledge for describing radiation and reflection of light. In Sect. 5 we describe experiments on (1) learning abstract definitions of polygons from artificial images, (2) predicting the light source direction and identification of ambiguities in images of the moon and microscopic images of illuminated micro-organisms and (3) identifying the ball in the RoboCup domain. Finally, we conclude and discuss further work in Sect. 6.

2 Related work

Statistical machine learning based on low-level feature extraction has been increasingly successful in image classification (Rautaray and Agrawal 2015). However, high-level vision, involving interpretation of objects and their relations in the external world, is still relatively poorly understood (Cox 2014). Since the 1990s *perception-by-induction* (Gregory 1998) has been the dominant model within computer vision, where human perception is viewed as inductive inference of hypotheses from sensory data. The idea originated in the work of the nineteenth century physiologist (von Helmholtz 1962). The approach described in this paper is in line with *perception-by-induction* in using ILP for generating high-level perceptual hypotheses by combining sensory data with a strong bias in the form of explicitly encoded

background knowledge. Whilst Gregory (1974) was one of the earliest to demonstrate the power of the Helmholtz's perception model for explaining human visual illusion, recent experiments (Heath and Ventura 2016) show Deep Neural Networks fail to reproduce human-like perception of illusion. This contrasts with results in Sect. 5.2, in which LV achieves analogous outcomes to human vision.

Early work in Computer Vision investigated the interaction between visual analysis, linguistic descriptions and geometric models (Waltz 1980; Huffman 1971). In some such approaches visual illusions were identified by testing logical models of images for contradictions (Barrow and Tenenbaum 1981). However, these techniques were based on preformulated models, and did not use machine learning augmented by background knowledge in the fashion described in this paper. Preformulated models are also used in more recent work to capture, for instance, the movement of a human being walking (Hogg 1983) or a hyperbolic curve involved in analysing images of penetrating radar (Olhoeft 2000). However, these techniques lack the flexibility of our Logical Vision approach to combine a set of primitive models in a modular fashion to form a set of composite structured and re-useable models from an image.

Shape-from-shading (Horn 1989; Zhang et al. 1999) is a key computer vision technology for estimating low-level surface orientation in images. Unlike our approach for identifying concavities and convexities, shape-from-shading generally requires observation of the same object under multiple lighting conditions. By using background knowledge as a bias we reduce the number of images for accurate perception of high-level shape properties such as the identification of convex and concave image areas.

ILP has previously been used for learning concepts from images. For instance, in Cohn et al. (2006) object recognition is carried out using existing low-level computer vision approaches, with ILP being used for learning general relational concepts from this already symbolised starting point. Farid and Sammut (2014a, b) adopted a similar approach, extracting planar surfaces from a 3D image of objects encountered by urban search and rescue robots and household objects, then using ILP to learn relational descriptions of those objects. By contrast, LV (Dai et al. 2015) uses ILP to provide a bridge from very low-level features, such as high contrast points, to high-level interpretation of objects. The present paper extends the earlier work on LV by implementing a noise-proofing technique, applicable to real images, and extending the use of generic background knowledge to allow the identification of objects, such as light sources, not directly identifiable within the image itself.

Various statistics-based techniques, making use of high-level vision, have been proposed for one- or even zero-shot learning (Palatucci et al. 2009; Vinyals et al. 2016). They usually start from an existing model pre-trained on a large corpus of instances, and then adapt the model to data with unseen concepts. Approaches can be separated into two categories. The first exploits a mapping from images to a set of semantic attributes, then high-level models are learned based on these attributes (Lampert et al. 2014; Mensink et al. 2011; Palatucci et al. 2009). The second approach uses statistics-based methods, pre-trained on a large corpus, to find localized attributes belonging to objects but not the entire image, and then exploits the semantic or spatial relationships between the attributes for scene understanding (Hu et al. 2016; Li et al. 2014; Duan et al. 2012). Unlike these approaches, we focus on one-shot from scratch, i.e. high-level vision based on just *very low-level features* such as high contrast points.

Machine learning is used extensively in robotics, mainly to learn perceptual and motor skills. Current approaches for learning perceptual tasks include Deep Learning and Convolutional Neural Networks (Krizhevsky et al. 2012; Redmon et al. 2016). The different approaches to vision in RoboCup can be seen in the SPQR team's use of convolutional neural networks (Suriani et al. 2016) and the *ad hoc*, but effective method used by the 2016 SPL

Table 1 Metarules used in this paper. Uppercase letters P, Q, R, S denote existentially quantified variables. Lowercase letters u, x, y , and z are universally quantified

Name	Metarule
PropObj1	$P(obj1) \leftarrow$
PropObj2	$P(obj2) \leftarrow$
PropLight	$P(light) \leftarrow$
Conjunct3	$P(x, y, z) \leftarrow Q(x, y, z), R(x, y, z)$
Chain3	$P(u, x, y) \leftarrow Q(u, x, z), R(u, z, y)$
Chain32	$P(u, x, y) \leftarrow Q(u, x, z), R(z, y)$
PrePost3	$P(x, y, z) \leftarrow Q(x, y), R(x), S(z)$
Pre2	$P(x) \leftarrow Q(x); R(x, y)$
Post2	$P(x, y) \leftarrow Q(x, y), R(y)$

champions, B-Human (Rofer et al. 2016). This approach clearly depends on domain knowledge that has been acquired by the human designers. However, the approach described in this paper promises the possibility that similar knowledge could be acquired through machine learning.

3 Framework

The framework for LV is a special case of MIL.

Meta-Interpretive Learning Given background knowledge B and examples E the aim of a MIL system is to learn a hypothesis H such that $B, H \models E$, where $B = B_p \cup M$, B_p is a set of Prolog definitions and M is a set of *metarules* (see Table 1). MIL (Muggleton et al. 2014b, a, 2015; Cropper and Muggleton 2015, 2016) is a form of ILP based on an adapted Prolog meta-interpreter. A standard Prolog meta-interpreter proves goals by repeatedly fetching first-order clauses whose heads unify with the goals. By contrast, a MIL learner proves the set of all examples by fetching higher-order metarules (Table 1) whose heads unify with the goals. The resulting meta-substitutions are saved, allowing them to be used to generate a hypothesised program which proves all the examples by substituting the meta-substitutions into corresponding metarules. Use of metarules and background knowledge helps minimise the number of clauses n of the minimal consistent hypothesis H and consequently the number of examples m required to achieve error below ϵ bound. Cropper and Muggleton (2016) shows n dominates the upper bound for m .²

Logical Vision In LV (Dai et al. 2015), the background knowledge B , in addition to Prolog definitions, contains a set of one or more named images I . The examples describe properties associated with I .

4 Implementation

4.1 Noise tolerant Meta-Interpretive Learning

The MIL framework described in the previous section has been implemented in a system called *Metagol* (Muggleton et al. 2014a, b, 2015; Cropper and Muggleton 2015, 2016).

² p predicates and M metarules $m \geq \frac{n \ln|M| + p \ln(3n) + \ln \frac{1}{\epsilon}}{\epsilon}$.

Algorithm 1: *Metagol_{NT}(B, E, v, n)*

```

Input : Background knowledge B; Set of (noisy) examples E; Parameter about noise level v and
         number of iterations n
Output: Hypothesis Hmax
1 max_score = 0;
2 max = 1;
3 for each i ∈ [1, n] do
   /* Randomly select examples from E w.r.t. noise level v           */
4   Tri = randSample(E, v);
   /* Leave the rest of examples for validation                       */
5   Tsi = E - Tri;
   /* Call Metagol and save the learned hypothesis in Hi           */
6   Hi = learn(B, Tri);
   /* Evaluate the learned hypothesis Hi on validating set         */
7   Ei = evaluate(B, Hi, Tsi);
8   if max_score < Ei then
9     | Hmax = Hi;
10    | max_score = Ei;
11  end
12 end
13 Return Hmax;

```

In this section we describe a noise tolerant version of *Metagol* called *Metagol_{NT}*.³ The standard *Metagol* implementation uses a modified Prolog meta-interpreter to backtrack through the space of hypotheses which prove all training examples. This strategy is consistent with an assumption of noise-free examples. Because of backtracking, standard methods for handling noise, such as accepting a user-defined maximum number of negative examples (used in the ILP systems Progol and Aleph), are inefficient for *Metagol*.⁴ For this reason, a more efficient noise-handling method is required.

The noise tolerant version of *Metagol* (i.e. *Metagol_{NT}*) used in this paper, finds hypotheses consistent with randomly selected subsets of the training examples and evaluates each resulting hypothesis on the remaining training set, returning the hypothesis with the highest score. The size of the training samples and the number of iterations (i.e. number of random samples) are user defined parameters. As shown in Algorithm 1, *Metagol_{NT}* is implemented as a wrapper around *Metagol* and returns the highest score hypothesis *H_{max}* learned from randomly sampled examples from *E* after *n* iterations. The sample size is controlled by $v = (k^+, k^-)$, where k^+ and k^- are the number of sampled positive and negative examples respectively, reflecting the noise level in the dataset.

4.2 Logical Vision

Our implementation of Logical Vision, called *LogVis*, is shown in Algorithm 2. The input consists of a set of images *I*, background knowledge *B* including both Prolog primitives *B_p* and metarules *M*, a set of training examples *E* of the target concept, *Metagol_{NT}*'s parameters *v* and *n*.

The procedure of *LogVis* is divided into two stages. The first stage is to extract symbolic background knowledge from images, which is done by the *visualAbduce* function. By

³ Available from https://github.com/atnezhad/Metagol_NT.

⁴ For example, a naive approach which ignores the label of the examples up to *k* times, has a time complexity of $O(\binom{m}{k})$, where *m* is the total number of examples.

Algorithm 2: *LogVis(I, B, E, v, n)*

```

Input : Training images  $I$ ; Background knowledge  $B$ ; Set of (noisy) examples  $E$ ; Parameter about
noise level  $v$  and number of iterations  $n$ .
Output: Hypothesised logic program  $H$ .
/* Initialise the knowledge base of visual primitives */
1  $B_v = \Phi$ ;
2 for each image  $i \in I$  do
3   /* Do visual abduction to get facts of visual primitives  $P$  */
4    $P_i = \text{visualAbduce}(i, B)$ ;
5    $B_v = B_v \cup P_i$ ;
6 end
/* Call  $\text{Metagol}_{NT}$  to learn a model */
7  $\text{Model} = \text{Metagol}_{NT}(B \cup B_v, E, v, n)$ ;
8 Return  $\text{Model}$ ;
    
```

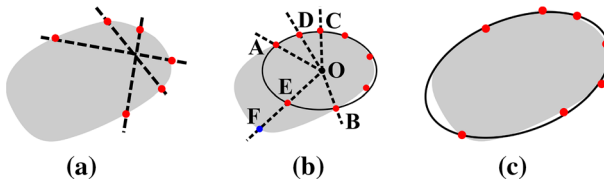


Fig. 3 Object detection: **a** sampled lines with edge points; **b** fitting of initial ellipse centred at O . Hypothesis tested using new edge points halfway between existing adjacent points. **c** Revised hypothesis tested until hypothesis passes test

including abductive theories in $B_p \in B$, *visualAbduce* can abduce. Points, lines, ellipses and even complex mid-level visual representations such as super-pixels (see Sect. 5.3). In our implementation, *visualAbduce* can take logic rules, statistical models and functions from a computer vision toolbox as background knowledge, which provide visual primitives. This makes *LogVis* flexible in learning many kinds of concepts. More details about visual abduction are introduced in Sect. 5.2.

The second stage of *LogVis* simply calls the noise-tolerant MIL system *Metagol_{NT}* to induce a hypothesis for the target concept, as both abduced visual primitives B_v and training examples E from an image dataset can be noisy.

Visual abduction The target of visual abduction is to obtain symbolic interpretation of images for further learning. The abduced logical facts are groundings of primitives defined in the background knowledge B_p . For example, in order to learn the concept of a polygon one at least needs to extract points and edges from an image. When the data is noise-free, this can be done by sampling high-contrast pixels from the image, such as the background knowledge about *edge_point* applied in Dai et al. (2015).

However, for real images that contain a degree of noise, we can include a statistical model in *visualAbduce* and use it to implement a noise-robust version of *edge_point*. For example, in the *Protist* and *Moon* experiments of Sect. 5, the *edge_point/l* predicate calls a pre-trained statistical image background model which can categorise pixels into foreground or background points using Gaussian models or image segmentation.

Furthermore, we can use an abductive theory about shapes to abduce objects. For example, in real images many objects of interest are composed of curves and can be approximated by ellipses or circles. Therefore we can include background knowledge about them in *visualAbduce* to perform ellipse and circle abduction, as shown in Fig. 3. The

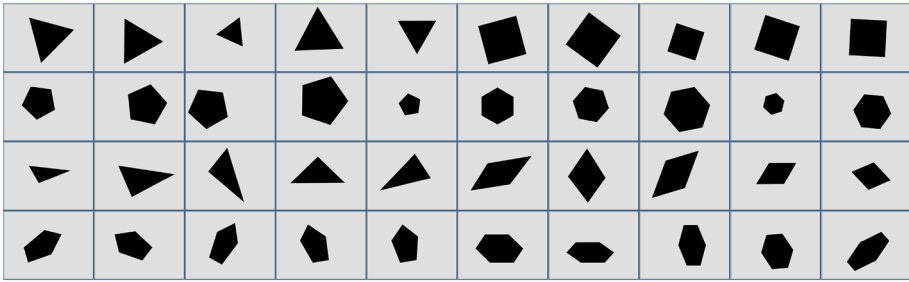


Fig. 4 Experiment 1—examples of the concept *regular polygon* used in Experiment 1. The first two rows are positive examples and the second two rows negative

abduced objects will take the form $elps(Centre, Parameter)$ or $circle(Centre, Radius)$ where $Centre = [X, Y]$ is the shape's centre, $Parameter = [A, B, Tilt]$ are the axis lengths and tilting angle and $Radius$ is the circle radius. The computational complexity of the abduction procedure is $O(rkn)$, where n is the number of `edge_points`, and k is the number of iteration of the ellipse fitting algorithm. r is the time required for resampling when the fitted object is not accurate enough, hence it is a constant that reflects the noise level of the input image.

In *LogVis*, background knowledge about visual primitives is implemented as logical predicates in a library, including basic geometrical concepts and extractors for low-level computer vision features such as the colour histogram and super-pixels. Users can implement their own background knowledge for visual abduction based on these primitives to address different kinds of problems flexibly.

5 Experiments

5.1 Experiment 1

In the first experiment [detailed report in Dai et al. (2015)] we compared a noise-free variant of the *LogVis* algorithm (referred to as LV_{Poly}) with statistics-based approaches on the task of learning simple geometrical concepts (see example images in Fig. 4).

Materials and methods We used Inkscape⁵ to randomly generate 3 labelled image datasets for 3 polygon shape learning tasks respectively. Training sets contain 40 examples. For simplicity, the images are binary-coloured, each image contains one polygon. Target concepts are: (1) `triangle/1`, `quadrangle/1`, `pentagon/1` and `hexagon/1`; (2) `regular_poly/1` (regular polygon); (3) `right_tri/1` (right triangle). All the datasets were partitioned into fivefold respectively, 4 of them were used for training and the remaining one is for testing, thus each experiment was conducted 5 times.⁶

Results and discussion Table 2 compares the predictive accuracies of an implementation of LV_{Poly} versus several statistics-based computer vision algorithms. We used a popular statistics-based computer vision toolbox VLFeat (Vedaldi and Fulkerson 2008) to implement the statistical learning algorithms. The experiments are carried with different kinds of features. Because the sizes of datasets are small, we used a support vector machine [libSVM (Chang

⁵ <http://inkscape.org>.

⁶ Code at <https://github.com/haldai/Logic-Vision>.

Table 2 Predictive accuracy of learning simple geometrical shapes from single object training sets of size 40

ACC	tri	pen	hex	reg	r_tri
HOG	0.83 ± 0.04	0.73 ± 0.03	0.75 ± 0.07	0.63 ± 0.08	0.74 ± 0.04
DSIFT	0.82 ± 0.05	0.64 ± 0.04	0.71 ± 0.03	0.71 ± 0.05	0.77 ± 0.07
LBP	0.87 ± 0.05	0.67 ± 0.03	0.73 ± 0.03	0.65 ± 0.05	0.75 ± 0.05
CNN	0.91 ± 0.01	0.75 ± 0.00	0.84 ± 0.02	0.59 ± 0.06	0.85 ± 0.04
C+d+L	0.82 ± 0.01	0.76 ± 0.01	0.76 ± 0.01	0.64 ± 0.05	0.80 ± 0.04
LV _{Poly}	1.00 ± 0.00	1.00 ± 0.00	0.99 ± 0.01	1.00 ± 0.00	1.00 ± 0.00

Maximal results are indicated in bold type

and Lin 2011]) as classifier. The parameters are selected by fivefold cross-validation. The features we have used in the experiments are as follows: **HOG**, Histogram of Oriented Gradients (Dalal and Triggs 2005), **Dense-SIFT**, Scale Invariant Feature Transform (Lowe 2004), **LBP**, Local Binary Pattern (Ojala et al. 2002), **CNN**, Convolutional Neural Network (CNN) (Simonyan and Zisserman 2015). We also compare with a combination of the above feature sets (i.e. **C+d+L**). According to Table 2, given 40 training examples the prediction accuracies for LV_{Poly} are significantly better than other approaches.

5.2 Experiment 2

This subsection describes experiments comparing one-shot LV with multi-shot statistics-based learning.⁷ In this experiments, we investigate the following null hypothesis:

Null hypothesis One-shot LV cannot learn models with accuracy comparable to thirty-shot statistics-based learning.

Materials We collected two real image datasets for the experiments: (1) **Protists** drawn from a microscope video of a *Protists* micro-organism, and (2) **Moons** a collection of images of the moon drawn from Google images. The instances in *Protists* are coloured images, while the images in *Moons* come from various sources and some of them are grey-scale. For the purpose of classification, we generated the two datasets by rotating images through 12 clock angles.⁸ Datasets consist of 30 images for each angle, providing a total of 360 images. Each image contains one of four labels as follows: *North* = {11, 12, 1} clocks, *East* = {2, 3, 4} clocks, *South* = {5, 6, 7} clocks, and *West* = {8, 9, 10} clocks. Examples of data and the labelling are shown in Fig. 5. As we can see from the figure, there is high variance in the image sizes and colours.

Methods The aim is to learn a model to predict the correct category of light source angle from real images. For each dataset, we randomly divided the 360 images into training and test sets, with 128 and 232 examples respectively. To evaluate the performance, the models were trained by randomly sampling 1, 2, 4, 8, 16, 32, 64 and 128 images from the training set. The sequences of training and test instances are shared by all compared methods. The random partition of data and learning are repeated 5 times.

Logical Vision In the experiments, we used the grey intensity of both image datasets for LV. The hyper-parameter T in Algorithm 2 is set at 11 by validating one-shot learned models on the rest of the training data. To handle image noise, we use a background model as the statistics-based estimator for predicate $edge_point/I$. When $edge_point([X,Y])$ is called, a

⁷ Data and code at <https://github.com/haldai/LogicalVision2>.

⁸ Clock face angle between 12 and each hour position in {1...12}.

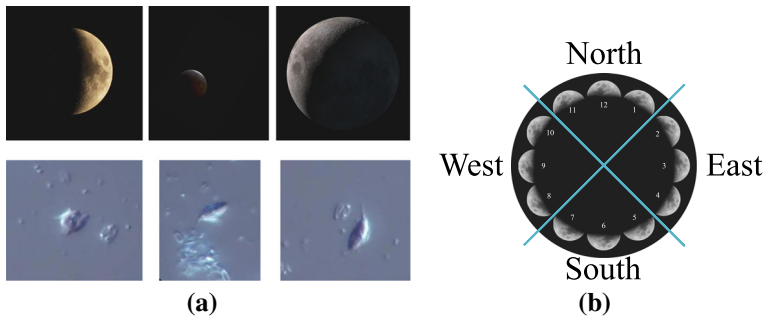


Fig. 5 Illustrations of *Moons* and *Protists* data: **a** examples of the datasets, **b** four classes for twelve light source positions

vector of colour distribution (which is represented by a histogram of grey-scale value) of the 10×10 region centered at (X, Y) is calculated, then the background model is applied to determine whether this vector represents an edge point. The parameter of neighborhood region size 10 is chosen as a compromise between accuracy and efficiency after having tested it ranging from 5 to 20. The background model is trained from five randomly sampled images in the training set by providing the bounding box of the objects.

Statistics-based Classification The experiments with statistics-based classification were conducted in different colour spaces combined with various features. Firstly, we performed feature extraction to transform images into fixed length vectors. Next SVMs [libSVM (Chang and Lin 2011)] with RBF kernel were applied to learn a multiclass-classifier model. Parameters of the SVM are chosen by cross validation on the training set. Like LV, we used grey intensity from both image datasets for the experiments. For the coloured *Protists* dataset, we transformed the images to **HSV** and **Lab** colour spaces to improve the performance. Since the image sizes in the dataset are irregular, during the object detection stage of LV, we used background models and computer graphics techniques (e.g. curve fitting) to extract the main objects and unified them into same sized patches for feature extraction. The sizes of object patches were 80×80 and 401×401 in *Protists* and *Moons* respectively. For the feature extraction process, we avoided descriptors which are insensitive to scale and rotation, instead we selected the luminance-sensitive features **HOG** and **LBP**. The Histogram of Oriented Gradient (HOG) (Dalal and Triggs 2005) is known for its ability to describe the local gradient orientation in an image, and widely used in computer vision and image processing for the purpose of object detection. Local binary pattern (LBP) (Ojala et al. 2002) is a powerful feature for texture classification by converting the local texture of an image into a binary number.

In the *Moons* task, LV and the compared statistics-based approach both used geometrical background knowledge for fitting circles (though in different forms) during object extraction. However, in the *Protists* task, the noise in images always caused poor performance in automatic object extraction for the statistics-based method. Therefore, we provided additional supervision to the statistics-based method consisting of bounding boxes labelling the position of the main objects in both training and test images during feature extraction. By comparison LV discovers the objects from raw images without any label information.

Results Figure 6a shows the results for *Moons*. Note that performance of the statistics-based approach only surpasses one-shot LV after 100 training examples. In this task, background knowledge involving circle fitting exploited by LV and statistics-based approaches are similar, though low-level features used by the statistics-based approach are first-order information (grey-scale gradients), which is stronger than the zeroth-order information (grey-scale value)

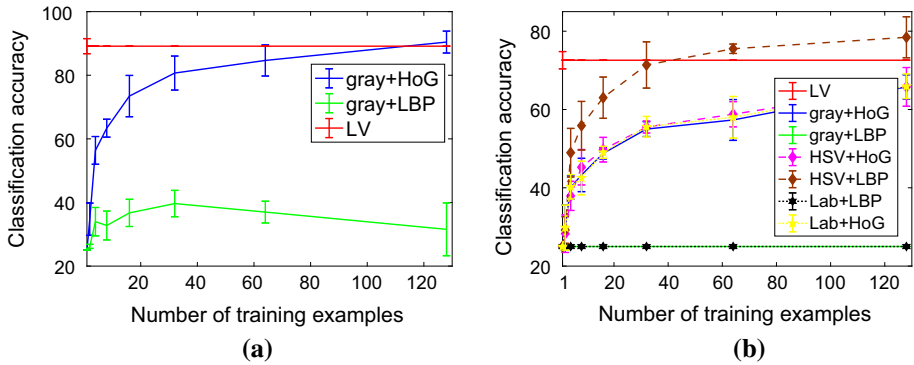


Fig. 6 Classification accuracy on the two datasets. **a** *Moons*. **b** *Protists*

```

clock_angle(A,B,C):-
    clock_angle1(A,B,D),
    light_source_angle(A,D,C).
clock_angle1(A,B,C):-
    highlight(A,B),
    clock_angle2(A),clock_angle3(C).
clock_angle2(obj1).
clock_angle3(light).
    
```

Fig. 7 Abductive program learned by LV: *clock_angle3* denotes the clock angle from *B* (highlight) to *A* (object). *high_light2* is a built-in predicate meaning *B* is the highlight part of *A*. *light_source_angle3* is an abducible predicate and the learning target. With background knowledge about lighting and compare this program with Fig. 8, we can interpret the invented predicate *clock_angle2* as *convex*, *clock_angle3* as *light_source_name*

used by LV. Results on *Protists* are shown in Fig. 6b. After 30+ training examples only one statistics-based approach outperforms one-shot LV. Since the statistics-based approaches have additional supervision (bounding box of main object) in the experiments, improved performance is unsurprising. The results of LV in Fig. 6a, b are represented by horizontal lines. When the number of training examples exceeds one, LV performs multiple one-shot learning and selects the most frequent output (see Algorithm 2), which we found is always in the same equivalent class in LV’s hypothesis space. This suggests LV learns the optimal model in its hypothesis space from a single example. The learned program is shown in Fig. 7.

The results in Fig. 6 demonstrate that *Logical Vision* can learn an accurate model using a single training example. By comparison, the statistics-based approaches require 40 or even 100 more training examples to reach similar accuracy, which refutes the null hypothesis. However, the performance of LV heavily relies on the accuracy of the statistical estimator of *edge_point/1*, because the mistakes of edge points detection will harm the shape fitting and consequently the accuracy of main object extraction. Unless we train a better *edge_point/1* classifier, the best performance of *LV* is limited as Fig. 6 shows.

LV is implemented in SWI-Prolog (Wielemaker et al. 2012) with multi-thread processing. Experiments were executed on a laptop with Intel i5-3210M CPU (2.50GHz), the time costs of object discovery are 9.5 and 6.4 s per image on *Protists* and *Moons* dataset respectively; the average running time *Metagol* procedure is 0.001 s on both datasets.

Protists and *Moons* contain only convex objects. If instead we provide images with concave objects (such as Fig. 9), LV learns a program such as Fig. 8. Here the invented predicate

Fig. 8 Program learned by LV when concave objects are given as training examples

```

clock_angle(A,B,C):-
  clock_angle1(A,B,D),clock_angle4(A,D,C).
clock_angle1(A,B,C):-
  highlight(A,B),clock_angle2(A),clock_angle3(C).
clock_angle4(A,B,C):-
  light_source_angle(A,B,D),opposite_angle(D,C).

```




Fig. 9 An image of a crater on Mars and the 180° rotated version. **a** Crater. **b** Flipped crater. Credit: NASA/JPL/University of Arizona

Fig. 10 Interpreted BK learned by LV

```

clock_angle(O,H,A):-
  highlight(O,H),convex(O),light_source(L),
  light_source_angle(O,L,A).
clock_angle(O,H,A):-
  highlight(O,H),concave(O),light_source(L),
  light_source_angle(O,L,A1),opposite(A1,A).

```

clock_angle2/1 can be interpreted as *concave* because its interpretation can be related to the appearance of *opposite_angle/2*.

Discussion: Learning ambiguity

Figure 9 shows two images of a crater on Mars, where Fig. 9b is a 180° rotated image of Fig. 9a. Human perception often confuses the convexity of the crater in such images.⁹ This phenomenon, called the *crater/mountain illusion*, occurs because human vision usually interprets pictures under the default assumption that the light is from the top of the image.

LV can use MIL to perform abductive learning. We show below that incorporation of generic recursive background knowledge concerning light enables LV to generate multiple mutually inconsistent perceptual hypotheses from real images. To the authors’ knowledge, such ambiguous prediction has not been demonstrated previously with machine learning.

Recall the learned programs from Figs. 7 and 8 from the previous experiments. If we rename the invented predicates we get the general theory about lighting and convexity shown in Fig. 10.

Now we can use the program as a part of interpreted background knowledge for LV to do abductive learning, where the abducible predicates and the rest of background knowledge are shown in Fig. 11.

⁹ <http://www.universetoday.com/118616/do-you-see-a-mountain-or-a-crater-in-this-picture/>.

Abducibles	
prim(convex/1).	prim(concave/1).
prim(light_source/1).	prim(light_source_angle/3).
Compiled BK	
% "obj1" is an object abducted from image, "obj2" is % the brighter part of "obj1"; "observer" is the camera	
contains(obj1,obj2).	brighter(obj2,obj1).
observer(observer).	reflector(obj2).
light_path(X,X).	
light_path(X,Y):-unobstructed(X,Z), light_path(Z,Y).	
Interpreted BK	
highlight(X,Y):-	
contains(X,Y),brighter(Y,X),light_source(L),	
light_path(L,R),reflector(R),light_path(R,O),	
observer(O).	

Fig. 11 Background knowledge for learning ambiguity from images

If we input Fig. 9a to LV, it will output four different abductive hypotheses for the image, as shown in Fig. 12.¹⁰ From the first two results we see that, by considering different possibilities of light source direction, LV can predict that the main object (which is the crater) is either convex or concave, which shows the power of learning ambiguity. The last two results are even more interesting: they suggest that *obj2* (the highlighted part of the crater) might be the light source as well, which indeed is possible, though seems unlikely.¹¹

5.3 Experiment 3

In this subsection we describe the experiments conducted on real images involving RoboCup¹² soccer where the task is to locate the football. We address this task in two stages: first we try to approximately locate the football in the image and then we use the model-driven technique of *Logical Vision* to abduce its location and shape. By doing this, one can estimate the size of the football, recognise occluded footballs and deduce depth information from the images.

Dataset and task The dataset contains 377 colour images sampled from a video of the robot's camera view of the football field. As Fig. 13 shows, the scene of this dataset contains the green field, a robot, and a football. The original size of the images are 480×720 . In this experiment they have been scaled into 240×360 for reducing the computational complexity.

This task is more difficult than those in the previous experiments. The objects in the images are more complex and contain more noise. Therefore it is difficult to learn a hypothesis using simple primitives such as "edge_point". For example, the robot and football contain many edges so the original line sampling based abduction used by *Logical Vision* will become a large-scale combinatorial optimisation problem. Moreover, in 41 of the images the football is either occluded by or connected to other objects, and in 40 images there is no football at all.

To address the challenges, we consider a two-staged learning procedure. The first sub-task is to quickly find candidate locations of the footballs, which can reduce the search space of

¹⁰ Code also at <https://github.com/haldai/LogicalVision2>.

¹¹ The result can be reproduced and visualized by the example in Logical Vision 2 repository.

¹² www.robocup.org.

Depiction	Hypothesis
<p>(a)</p>	<pre>light_source(light). light_source_angle(obj1,light,south). convex(obj1).</pre>
<p>(b)</p>	<pre>light_source(light). light_source_angle(obj1,light,north). concave(obj1).</pre>
<p>(c)</p>	<pre>light_source(obj2). light_source_angle(obj1,obj2,south). convex(obj1).</pre>
<p>(d)</p>	<pre>light_source(obj2). light_source_angle(obj1,obj2,north). concave(obj1).</pre>

Fig. 12 Depiction of abducted hypotheses from Fig. 9a

the fine grained football discovery. The second sub-task is to use *Logical Vision* to abduce the location and shape of the football from the candidate positions.

For the first sub-task, we use a super-pixel algorithm (Achanta et al. 2012) to segment the images into small regions, which can serve as primitives for estimating the location of football. Super-pixel algorithms are able to group pixels into atomic regions that capture image redundancy, greatly reducing the complexity of subsequent image processing tasks. The super-pixel algorithm implementations we used are `OpenCV_contrib`¹³ (Bradski 2000).

¹³ https://github.com/opencv/opencv_contrib.




Fig. 13 Examples of football images: **a** the football is clearly separated from other objects, **b** part of the football is located outside of the image, **c** the football is occluded by the robot

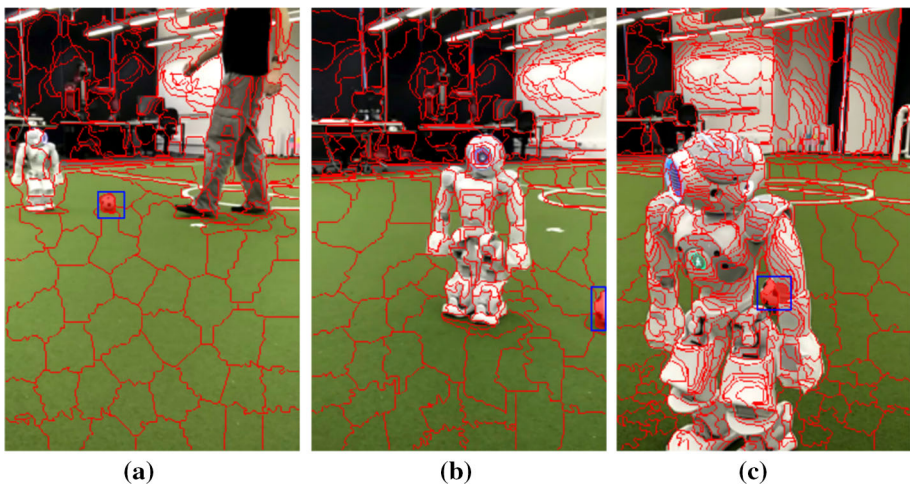


Fig. 14 Super-pixel segmented data of the images in Fig. 13, where the blue boxes are the original bounding boxes of the images, the super-pixels filled with red colour are the positive super-pixels according to the bounding boxes. Note that in **(b)** and **(c)**, although the footballs have been split into multiple super-pixels, they are all labelled as positive examples (Color figure online)

The tuned parameter is the size of each super-pixel, which ranges from 10 to 30 with step size 5. During data transformation, we use the football bounding boxes shipped with original images to label the super-pixels: those which have 95% area inside of a bounding box of footballs (which is the label information in original data) are labelled as positive examples with predicate “ball_sp”. The rest are labelled as negatives. Examples from the dataset are shown in Fig. 14. The second sub-task, model-driven football abduction, directly takes “ball_sp” and an abductive theory as input and outputs the circle parameters (centre and radius), where “ball_sp” should be the result produced by the classification model learned in the first stage.

Experiment: Football super-pixel classification This experiment is related to the first sub-task described above, i.e. locating the football from super-pixel segmented images. In this experiment we compare the performance of *Metagol_{NT}* versus a statistical learner [we choose the CART algorithm (Breiman et al. 1984)¹⁴] and investigate the same null hypothesis used in Sect. 5.2.

Materials and methods In this experiment we use the super-pixel dataset as described above. Each super-pixel is regarded as a symbolic object in the background knowledge. We extract some basic properties, such as size, location and colour distribution as features. The colour distribution is represented by the proportion of white, grey, black and green pixels inside a super-pixel, which is identified by **Lab** values of the pixels. Moreover, we exploit the neighbourhood relationship between super-pixels, which is represented by the “next_to/2” predicate.¹⁵

In this experiment we randomly sample 128 images for the training and the remaining 249 images for testing. Similar to the *Protists* and *Moons* experiments in Sect. 5.2, we randomly sample 1, 2, 4, 8, 16, 32, 64, 128 images from the training set for learning the classification model. Random data partitioning is performed 5 times. The positive training examples (both for the statistical learner and the relational learner) are football super-pixels from each of 1, 2, 4, 8, 16, 32, 64, 128 images and the same number of negative examples (i.e. non-football super-pixels) are randomly sampled from the same set of training images. Similarly, for the test data the negative examples are randomly sampled from non-football super-pixels in the test images. For relational learning (i.e. *Metagol_{NT}*), background predicates *mostly_white/1*, *partly_white/1*, *mostly_black/1*, *partly_black/1*, etc were defined based on the colour distribution of super-pixels. For example the following background definitions describe a super-pixel which is mostly white or partly white:

```
mostly_white(S):- white(S, P), P > 0.6.
partly_white(S):- white(S, P), P > 0.4, P <= 0.6.
```

The background knowledge for the relational learner also includes the neighbourhood relationship between super-pixels, i.e. “next_to/2” predicates.

In this experiment the following parameters were used for the relational learner, i.e. *Metagol_{NT}(B, E, v, n)* in Algorithm 1. In addition to the above mentioned background knowledge, *B* includes the *Pre2* and *Post2* Meta-rules from Table 1.

E is the set of positive and negative training examples as described above. The size of randomly selected training examples $Tr_i \subset E$ in each iteration *i* of Algorithm 1 and the number of iterations *n* can be set according to the expected degree of noise. Given that the expected error rate in the training data is not known in this problem, we choose an extreme case where Tr_i contains one randomly selected positive example (and one or two randomly selected negative examples) in different experiments. The number of iterations *n* was set to the number of positive examples in *E*.

For the statistics-based learner we use the CART decision tree algorithm (Breiman et al. 1984). The goal is to create a model that predicts the value of a target variable based on splitting the feature space. We choose CART as the compared method because we want to ensure the statistical model uses the same features as the relational model. Since the number of features, i.e. the green/white/grey/black pixel proportions, is relatively small, it is

¹⁴ CART was chosen since it is efficient and provides human-comprehensible output comparable to logic programs and execution of decision trees within the Robocup environment is sufficiently efficient (under 1/30th of a second) for localisation and decision making.

¹⁵ Dataset located at <https://github.com/haldai/LogicalVision2/tree/master-2.1/data>.

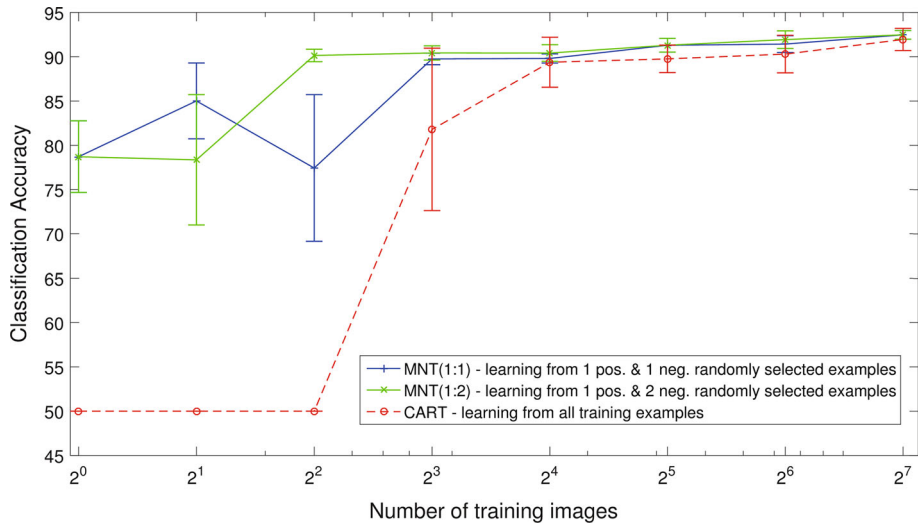


Fig. 15 Accuracy of $Metagol_{NT}$ vs CART in the task of football super-pixel classification

natural to choose a decision tree as the statistical learner. The maximum number of splits, is automatically selected by fivefold cross validation on the training data.

A second reason for the choice of decision trees is efficiency of execution. The robots in RoboCup soccer must operate in real-time, which means that all vision, localisation, decision making, localisation and locomotion tasks must be completed in the time it takes to capture the next camera frame, typically 1/30th of a second. Thus, the classifier in the vision system must be extremely efficient to execute. A decision tree, with only a few comparisons leading to a decision in the leaf node, satisfies these stringent timing requirements.

Results Figure 15 compares the predictive accuracy of the relational learner ($Metagol_{NT}$) vs the statistics-base learner (CART). As shown in the figure, $Metagol_{NT}$ achieves consistently higher accuracy than CART with the accuracy difference particularly high for small numbers of training examples. An example of the hypotheses found by the relational learner is as follows:

```
ball_sp(A):- partly_white(A), ball_sp_1(A, B).
ball_sp_1(A,B):- next_to(A, B), mostly_green(B).
```

Model-driven football abduction After narrowing down the candidate location of the football, *Logical Vision* is able to exploit geometrical background knowledge to perform model-driven abduction of the football's exact shape and position (i.e. its centre and radius as a circle). This is important in robotic football games since the robot can use this information to infer the distance between itself and the football. More importantly, by modelling the football with a circle, the robot can figure out the occlusion of the football by other robots and choose appropriate actions accordingly. We apply *Logical Vision* with an abductive theory for this task, whose abducible is "football/3". To sample edge points, *Logical Vision* draws random straight lines inside a super-pixel and its neighbourhood to return the points associated with a colour transition. Examples of football abduction are shown in Fig. 16.




Fig. 16 Ball abduction results of the images in Fig. 13. The blue points are the “edge_points” sampled by *Logical Vision*, the red curves are the abducted circles (Color figure online)

6 Conclusions and further work

Human beings often learn visual concepts from single image presentations (so-called one-shot-learning) (Lake et al. 2011). This phenomenon is hard to explain from a standard Machine Learning perspective, given that it is unclear how to estimate any statistical parameter from a single randomly selected instance drawn from an unknown distribution. In this paper we show that learnable generic logical background knowledge can be used to generate high-accuracy logical hypotheses from single examples. This compares with similar demonstrations concerning one-shot MIL on string transformations (Lin et al. 2014) as well as previous concept learning in artificial images (Dai et al. 2015). The experiments in Sect. 5 show that the LV system can accurately identify the position of a light source from a single real image, in a way analogous to scientists such as Galileo, observing the moon for the first time through a telescope or Hook observing micro-organisms for the first time through a microscope. In Sect. 5.2 we show that logical theories learned by LV from labelled images can also be used to predict concavity and convexity predicated on the assumed position of a light source. Section 5.3 shows how LV can be used effectively in real-time robot vision. Ball recognition in robot soccer is challenging because the ball is frequently occluded by other robots and the similarity in colours of the ball, robots and field lines makes the ball difficult to distinguish.

We have studied LV’s failure cases carefully. The main reason causing misclassification is the noise in images. The noise can cause misclassifications of `edge_point/1` since it is implemented with statistical models. The mistakes of `edge_point` detection will further affect the edge detection and shape fitting. As a result, the accuracy of the main object extraction is limited by both the noise level in input images and the power of statistical model of `edge_point/1`. Therefore, LV will fail too since the wrongly extracted objects are its inputs. However, if we train stronger models for detecting `edge_points`, the accuracy of LV will not increase either.

In further work we aim to investigate broader sets of visual phenomena which can naturally be treated using background knowledge. For instance, the effects of object obscuration; the interpretation of shadows in an image to infer the existence of out-of-frame objects; the

existence of unseen objects reflected in a mirror found within the image. All these phenomena could possibly be considered in a general way from the point of view of a logical theory describing reflection and absorption of light, where each image pixel is used as evidence of photons arriving at the image plane. In this further work we aim to compare our approach once more against a wider variety of competing methods.

The authors believe that LV has long-term potential as an AI technology with the potential for unifying the disparate areas of logical based learning with visual perception.

Funding Funding was provided by Royal Academy of Engineering (Grant No. NRC/P1415/133).

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Achanta, R., Shaji, A., Smith, K., Lucchi, A., Fua, P., & Susstrunk, S. (2012). SLIC superpixels compared to state-of-the-art superpixel methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(11), 2274–2282.
- Barrow, H. G., & Tenenbaum, J. M. (1981). Interpreting line drawings as three-dimensional surfaces. *Artificial Intelligence*, 17, 75–116.
- Bradski, G. (2000). *OpenCV library*. <http://opencv.org/>.
- Breiman, L., Friedman, J. H., Olshen, R. A., & Stone, C. J. (1984). *Classification and regression trees*. Belmont: Wadsworth.
- Chang, C.-C., & Lin, C.-J. (2011). LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2, 27:1–27:27.
- Cohn, A. G., Hogg, D. C., Bennett, B., Galata, A., Magee, D. R., & Santos, P. (2006). In: *Cognitive vision: Integrating symbolic qualitative representations with computer vision* (pp. 221–246). Berlin: Springer.
- Cox, D. (2014). Do we understand high-level vision? *Current Opinion in Neurobiology*, 25, 187–193.
- Cropper, A., & Muggleton, S. H. (2015). Logical minimisation of meta-rules within meta-interpretive learning. In *Proceedings of the 24th international conference on inductive logic programming* (pp. 65–78). LNAI 9046. Springer.
- Cropper, A., & Muggleton, S. H. (2016). Learning higher-order logic programs through abstraction and invention. In *Proceedings of the 25th international joint conference artificial intelligence (IJCAI 2016)* (pp. 1418–1424). IJCAI.
- Dai, W.-Z., Muggleton, S. H., & Zhou, Z.-H. (2015). Logical vision: Meta-interpretive learning for simple geometrical concepts. In *Late breaking paper proceedings of the 25th international conference on inductive logic programming* (pp. 1–16). CEUR.
- Dalal, N., & Triggs, B. (2005). Histograms of oriented gradients for human detection. In: *Proceedings of the 13rd IEEE computer society conference on computer vision and pattern recognition*, San Diego, CA (pp. 886–893). IEEE Computer Society.
- Duan, K., Parikh, D., Crandall, D. J., & Grauman, K. (2012). Discovering localized attributes for fine-grained recognition. In *Proceedings of the 25th IEEE conference on computer vision and pattern recognition*, Providence, RI (pp. 3474–3481). IEEE Computer Society.
- Farid, R., & Sammut, C. (2014a). Plane-based object categorisation using relational learning. *Machine Learning*, 94, 3–23.
- Farid, R., & Sammut, C. (2014b). Region-based object categorisation using relational learning. In *PRICAI 2014: Trends in artificial intelligence* (Vol. 8862, pp. 1106–1114). LNAI. Springer.
- Galilei, G. (2004). *The Herald of the Stars*. 1610. English translation by Edward Stafford Carlos, Rivingtons, London, 1880; edited by Peter Barker, Byzantium Press.
- Gregory, R. L. (1974). *Concepts and mechanics of perception*. London: Duckworth.
- Gregory, R. L. (1998). *Eye and brain: The psychology of seeing*. Oxford: Oxford University Press.
- Heath, D., & Ventura, D. (2016). Before a computer can draw, it must first learn to see. In: *Proceedings of the 7th international conference on computational creativity* (pp. 172–179).

- Hogg, D. (1983). Model-based vision: A program to see a walking person. *Image and Vision Computing*, 1, 5–20.
- Horn, B. K. P. (1989). *Obtaining shape from shading information*. Cambridge, MA: MIT Press.
- Hu, R., Xu, H., Rohrbach, M., Feng, J., Saenko, K., & Darrell, T. (2016). Natural language object retrieval. In *Proceedings of the 29th IEEE conference on computer vision and pattern recognition*, Las Vegas, NV (pp. 4555–4564). IEEE Computer Society.
- Huffman, D. A. (1971). Impossible objects as nonsense sentences. In B. Meltzer & D. Michie (Eds.), *Machine intelligence* (Vol. 6, pp. 295–323). Edinburgh: Edinburgh University Press.
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems*, 25, 1106–1114.
- Lake, B. M., Salakhutdinov, R., Gross, J., & Tenenbaum, J. B. (2011). One shot learning of simple visual concepts. In: *Proceedings of the 33rd annual conference of the cognitive science society* (pp. 2568–2573).
- Lampert, C. H., Nickisch, H., & Harmeling, S. (2014). Attribute-based classification for zero-shot visual object categorization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(3), 453–465.
- Li, Z., Gavves, E., Mensink, T., & Snoek, C. G. M. (2014). Attributes make sense on segmented objects. In: *Proceedings of 13th European conference on computer vision part IV*, Zurich, Switzerland (pp. 350–365). Springer.
- Lin, D., Dechter, E., Ellis, K., Tenenbaum, J. B., & Muggleton, S. H. (2014). Bias reformulation for one-shot function induction. In *Proceedings of the 23rd European conference on artificial intelligence (ECAI 2014)*, Amsterdam (pp. 525–530). IOS Press.
- Lowe, D. G. (2004). Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2), 91–110.
- Mensink, T., Verbeek, J. J. & Csurka, G. (2011). Learning structured prediction models for interactive image labeling. In *The 24th IEEE conference on computer vision and pattern recognition*, Colorado Springs, CO (pp. 833–840). IEEE Computer Society.
- Muggleton, S. H., Lin, D., Chen, J., & Tamaddoni-Nezhad, A. (2014a). Metabayes: Bayesian meta-interpretive learning using higher-order stochastic refinement. In G. Zaverucha, V. Santos Costa, & A. Marins Paes (Eds.), *Proceedings of the 23rd international conference on inductive logic programming (ILP 2013)*, Berlin. LNAI 8812 (pp. 1–17). Springer.
- Muggleton, S. H., Lin, D., Pahlavi, N., & Tamaddoni-Nezhad, A. (2014b). Meta-interpretive learning: Application to grammatical inference. *Machine Learning*, 94, 25–49.
- Muggleton, S. H., Lin, D., & Tamaddoni-Nezhad, A. (2015). Meta-interpretive learning of higher-order dyadic datalog: Predicate invention revisited. *Machine Learning*, 100(1), 49–73.
- Muggleton, S. H., De Raedt, L., Poole, D., Bratko, I., Flach, P., & Inoue, K. (2011). ILP turns 20: Biography and future challenges. *Machine Learning*, 86(1), 3–23.
- Ojala, T., Pietikainen, M., & Mäenpää, T. (2002). Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(7), 971–987.
- Olhoeft, G. R. (2000). Maximizing the information return from ground penetrating radar. *Journal of Applied Geophysics*, 43, 175–187.
- Palatucci, M., Pomerleau, D., Hinton, G., & Mitchell, T. M. (2009). Zero-shot learning with semantic output codes. In: *Advances in neural information processing systems* (Vol. 22, pp. 1410–1418). Curran Associates Inc.
- Rautaray, S. S., & Agrawal, A. (2015). Vision based hand gesture recognition for human computer interaction: A survey. *Artificial Intelligence Review*, 43, 1–54.
- Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2016). You only look once: Unified, real-time object detection. In: *The IEEE conference on computer vision and pattern recognition (CVPR)* (pp. 779–788).
- Rofer, T., Laue, T., Richter-Klug, J., & Thielke, F. (2016). B-human team description for RoboCup 2016, 2016. http://www.robocup2016.org/media/symposium/Team-Description-Papers/StandardPlatform/RoboCup_2016_SPL_TDP_B-Human.pdf.
- Simonyan, K., & Zisserman, A. (2015). Very deep convolutional networks for large-scale image recognition. In: *Proceedings of the 3rd international conference on learning representations*, San Diego, CA.
- Suriani, V., Albani, D., Youssef, A., Del Duchetto, F., Nania, A., Bloisi, D.D., et al. (2016). SPQR RoboCup 2016 standard platform league team description paper, 2016. http://www.robocup2016.org/media/symposium/Team-Description-Papers/StandardPlatform/RoboCup_2016_SPL_TDP_SPQR.pdf.
- Vedaldi, A., & Fulkerson, B. (2008). VLFeat: An open and portable library of computer vision algorithms. <http://www.vlfeat.org/>.

- Vinyals, O., Blundell, C., Lillicrap, T. P., Kavukcuoglu, K., & Wierstra, D. (2016). Matching networks for one shot learning. In: *Advances in Neural Information Processing Systems* (Vol. 29, pp. 3630–3638). MIT Press.
- von Helmholtz, H. (1962). *Treatise on physiological optics* (Vol. 3). New York: Dover Publications. (**Originally published in German in 1825**).
- Waltz, D. L. (1980). Understanding scene descriptions as event simulations. In: *Proceedings of the 18th annual meeting on association for computational linguistics* (pp. 7–11). Association for Computational Linguistics.
- Wielemaker, J., Schrijvers, T., Triska, M., & Lager, T. (2012). SWI-Prolog. *Theory and Practice of Logic Programming*, 12(1–2), 67–96.
- Zhang, R., Tai, P. S., Cryer, J. E., & Shah, M. (1999). Shape-from-shading: a survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(8), 670–706.