



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Human-machine scientific discovery

Citation for published version:

Tamaddoni-Nezhad, A, Bohan, DA, Milani, GA, Raybould, A & Muggleton, S 2021, Human-machine scientific discovery. in S Muggleton & N Chater (eds), *Human-Like Machine Intelligence*. Oxford University Press, pp. 279-315. <https://doi.org/10.1093/oso/9780198862536.003.0015>

Digital Object Identifier (DOI):

[10.1093/oso/9780198862536.003.0015](https://doi.org/10.1093/oso/9780198862536.003.0015)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Early version, also known as pre-print

Published In:

Human-Like Machine Intelligence

Publisher Rights Statement:

This is a draft of a chapter that has been accepted for publication by Oxford University Press in the book *Human-Like Machine Intelligence* edited by Stephen Muggleton and Nicholas Chater published in 2021.

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Human-Machine Scientific Discovery

Alireza Tamaddoni-Nezhad^{1,2}, David A. Bohan³, Ghazal Afroozi Milani², Alan Raybould⁴ and Stephen Muggleton²

¹University of Surrey, ²Imperial College London, UK, ³INRA, France and ⁴University of Edinburgh, UK

Abstract

Humanity is facing existential, societal challenges related to food security, ecosystem conservation, antimicrobial resistance, etc, and Artificial Intelligence (AI) is already playing an important role in tackling these new challenges. Most current AI approaches are limited when it comes to 'knowledge transfer' with humans, i.e. it is difficult to incorporate existing human knowledge and also the output knowledge is not human comprehensible. In this chapter we demonstrate how a combination of comprehensible machine learning, text-mining and domain knowledge could enhance human-machine collaboration for the purpose of automated scientific discovery where humans and computers jointly develop and evaluate scientific theories. As a case study, we describe a combination of logic-based machine learning (which included human-encoded ecological background knowledge) and text-mining from scientific publications (to verify machine-learned hypotheses) for the purpose of automated discovery of ecological interaction networks (food-webs) to detect change in agricultural ecosystems using the Farm Scale Evaluations (FSEs) of genetically modified herbicide-tolerant (GMHT) crops dataset. The results included novel food-web hypotheses, some confirmed by subsequent experimental studies (e.g. DNA analysis) and published in scientific journals. These machine-learned food-webs were also used as the basis of a recent study revealing resilience of agro-ecosystems to changes in farming management using GMHT crops.

1. Introduction

Humanity is facing existential, societal challenges related to the well-being and sustaining a growing population of 7.7 billion people and issues such as food security, the use of biotechnology in agriculture and medicine, antimicrobial resistance (AMR) and the emergence of new pathogens and pandemic diseases.

Scientists are today equipped with an ever-growing amount of human knowledge and empirical data in addition to advanced technologies such as Artificial Intelligence (AI). AI and Machine Learning are already playing an important role in tackling these new scientific challenges. For example, AI in the form of Deep Learning has recently been used in the discovery of a new candidate antibiotic which has been successfully tested against a range of antibiotic-resistant strains of bacteria (Stokes et. al, 2020).




Figure 1 Perspective of Human-Machine knowledge transfer as variants of AI research. a) Expert Systems (1980s) with a dependence on manual encoding of human knowledge, b) Deep Learning and Big Data in which humans are excluded from the encoded knowledge and c) Human-Like Computing (HLC) in which Humans and Computers jointly develop and share knowledge.

Despite great potential for new scientific discoveries, most current AI approaches, including Deep Learning, are limited when it comes to ‘knowledge transfer’ with humans. It is difficult to incorporate existing human knowledge and the output knowledge is not human comprehensible. Knowledge transfer is, however, a critically important part of Human-Machine Discovery which is necessary for collaboration between humans and AI.

Human-machine knowledge transfer is the subject of Human-Like Computing, also known as Third Wave of AI. Human-Like Computing (HLC) research aims to endow machines with human-like perception, reasoning and learning abilities which support collaboration and communication with human beings. Such abilities should support computers in interpreting the aims and intentions of humans based on learning and accumulated background knowledge.

Figure 1 shows the change in perspective which HLC represents in AI research, in particular with regards to knowledge transfer with humans. The idea of incorporating human knowledge in AI is not new and it was the basis of Expert Systems in 1980s where machines were dependent on being fed explicit knowledge from human experts (Figure 1a). However, incorporating existing knowledge and knowledge transfer are limited in the present black-box forms of AI where computers learn from Big Data, while humans are excluded from both the knowledge development cycle and the understanding of output knowledge (Figure 1b). In HLC, by contrast, humans and machines are viewed as co-developers of knowledge (Figure 1c). In the HLC world, we envisage a symmetric form of learning in which humans derive explicit knowledge from machines, and machines learn from humans and other data sources.

This form of two-way human-machine learning is also related to *ultra-strong* machine learning as defined by Michie (1988). Michie’s aim was to provide operational criteria for various qualities of machine learning which include not only predictive performance but also comprehensibility of learned knowledge. His weak criterion identifies the case in which the machine learner produces improved predictive performance with increasing amounts of data. The strong criterion additionally requires the learning system to provide its hypotheses in symbolic form. Lastly, the ultra-strong criterion extends the strong criterion by

requiring the learner to teach the hypothesis to a human, whose performance is consequently increased to a level beyond that of the human studying the training data alone.

In this chapter we demonstrate how a logic-based machine learning approach could meet the ultra-strong criterion and how a combination of this machine learning approach, text-mining and domain knowledge could enhance human-machine collaboration for the purpose of automated scientific discovery where humans and computers jointly develop and evaluate scientific theories.

As a case study, we describe a combination of the logic-based machine learning (which included human-encoded ecological background knowledge) and text-mining from scientific publications (to evaluate machine-learned hypotheses and also to identify potential novel hypotheses) for the purpose of automated discovery of ecological interaction networks (food-webs) from a large-scale agricultural dataset. Many of the learned trophic links were corroborated by the literature: in particular, links ascribed with high probability by machine learning corresponded with those having multiple references in the literature. In some cases, previously unobserved but high probability links were suggested and subsequently confirmed by subsequent experimental studies.

These machine-learned food-webs were also the basis of a recent study (Ma et. al., 2019) revealing resilience of agro-ecosystems to changes in farming management using genetically modified herbicide-tolerant (GMHT) crops.

This chapter is organised as follows. Sec. 2 describes the scientific problem and dataset. The knowledge gap for modelling agro-ecosystems is discussed in Sec. 3. Sec. 4 describes a machine learning approach for automated discovery of ecological networks. The ecological evaluation of the results and subsequent discoveries are discussed in Sec. 5. Sec. 6 concludes the chapter.

2. Scientific problem and dataset: Farm Scale Evaluations (FSEs) of GMHT crops

Humanity is facing great challenges to feed the growing population of 7.7 billion people and sustainable management of ecosystems and growth in agricultural productivity is at the heart of the United Nations' Sustainable Development Goals for 2030. Innovative agricultural management will be required to minimise greenhouse gas emissions and enrich biodiversity, provide sufficient nutritious food and maintain farmers' livelihoods and thriving rural economies. Predicting system-level effects will be crucial to introducing management that optimises delivery of many potentially conflicting objectives of agricultural, environmental and social policy.

Replacing existing conventional weed management with GMHT crops, for example, might reduce herbicide applications and increase crop yields. However, this requires an evaluation of the risks and opportunities owing to concerns about potential adverse impacts of GMHT crop management on biodiversity and the functioning of the agro-ecosystems.




Figure 2 Map of study fields in the FSEs. The circles show the locations of the field sites of spring-sown beet, maize and oilseed rape, and winter-sown oilseed rape overlain across UK.

The Farm Scale Evaluations (FSE) was a three-year study to test the effects of GMHT crop management on farmland biodiversity across the United Kingdom, and the details of farmland selection and crop field design are described in (Champion et al. 2003; Bohan et al. 2005). To summarize, a split-field design was used in 64 beet, 57 maize, 65 spring-sown oilseed rape and 65 winter-sown oilseed rape sites in the UK (see Figure 2). Each crop field was split approximately in half and a conventional and GMHT variety of one of the crops assigned randomly to each half. Plant and invertebrate species were sampled using a variety of standard ecological protocols. Taxa identity and abundance information were recorded within the field across all the sites. Approximately 60,000 field visits were made, sampling some 930,000 plants and 650,000 seeds that were identified to species. In excess of 2 million invertebrates were sampled, and 24,000 bees and 18,000 butterflies counted on the transect walks.

The overarching null hypothesis for the FSEs was that “there was no effect of the herbicide management of GMHT crops on biodiversity”, but with the expectation that effects on biodiversity would be mediated by a combination of the direct effects of herbicides killing weed plants and indirect effects on wider biodiversity through the loss of refuge and food resources provided by these weeds. The FSE scientists and steering committee agreed that a biologically significant effect on any taxon was a change in amount (count, density, biomass) of 50%, either up or down.

The sample data were analysed on a taxon by taxon basis using statistical approaches such as ANOVA (Perry et al., 2003). The null hypothesis was tested with a paired randomization test using the treatment effect, d (computed as $d = \log_{10}(\text{GM} + 1) - \log_{10}(\text{C} + 1)$), for the difference in count for a taxon due to management in the GM and Conventional half-fields. The results of the analyses demonstrated that there were significant changes to the amounts of some taxa of weeds, surface dwelling invertebrates and bees and butterflies in the different crops, with some going up and others down in the GMHT.

Assessment of the probable changes to biodiversity from adopting GMHT crops was used to inform decision-making by regulatory authorities and companies. For a variety of environmental policy and commercial reasons, none of the crops were commercialised in the UK. Nevertheless, the FSE dataset is the largest agro-ecological census dataset collected to date and it provided the agricultural Big Data used in the human-machine discovery of agro-ecological networks described in this chapter. Network reconstruction was done from invertebrate abundances sampled in the Vortis suction sampling and Pitfall trapping protocols. A new study using these machine-learned food-webs has also revealed that network-level responses in GMHT crop fields are remarkably similar in their composition, network properties, and responses to simulated trajectories of species removals, to their conventional counterparts, suggesting the resilience of agro-ecosystems to changes in farming management using GMHT crops (see Section 5).

3. The knowledge gap for modelling agro-ecosystems: ecological networks

The agro-ecological, mechanistic underpinning of ecosystem services, their response to change and how they interact is still poorly understood, as exemplified by the so-called “optimist’s scenario” (Pocock et al. 2012), which may be summarized as ‘the management of one ecosystem service, for improved outcomes, benefits the outcomes of all ecosystem services.’ The specific dependencies of one service on any other are only poorly understood and the validity of this scenario at system-relevant scales can only be guessed.

Since ecosystems are structured by flows of energy (biomass) between primary producer plants (autotrophs) and consumers (heterotrophs), such as invertebrates, mammals and birds (Lindeman 1942, Dickinson and Murphy 1998), food-webs are key explanations of ecosystem structure and dynamics that could be used to understand and predict responses to environmental change (Odum, 1974; Caron-Lormier et al., 2009; Cohen, Schittler, Raffaelli and Reuman, 2009, Woodward et al., 2012).

Still relatively few ecosystems have been described and detailed using food webs because establishing interactions, such as predation, between the many hundreds of species in an ecosystem is resource intensive, requiring considerable investment in field observation and laboratory experimentation (Ings et al 2009). Across such large datasets, it is often difficult to relate observational data sampled in protocols that have different basic metrics, such as density or activity-density or absolute abundance. Increasing the efficiency of testing for trophic links by filtering out unlikely interactions is typically not possible because of uncertainty about basic background knowledge of the network, such as whether any two species are likely even to come into contact and then interact (Ings et al 2009). In addition, it may require considerable analysis and interpretation to translate from the ecological ‘language’ of sample data (count, abundance, density, etc.) to the network language of nodes and links within a trophic network. Consequently, of those ecosystems that have been studied using trophic network approaches, component communities that provide known, valuable ecosystem services or those that are experimentally tractable or under threat have most often been evaluated (Ings et al 2009).

To make good decisions about ecosystem management, e.g. the management of agricultural land for the optimal delivery of ecosystem services, it is necessary to have theories that predict the effects of perturbation on ecosystems. Network ecology, and in particular food-webs, holds great promise as an approach to modeling and predicting the effects of perturbation on ecosystems. Networks of trophic links (i.e. food-webs) that describe the flow of energy/biomass between species, are important for making predictions about ecosystem structure and dynamics. However, relatively few ecosystems have been studied through detailed food-webs because establishing predation relationships between the many hundreds of species in an ecosystem is expensive, and in many cases, impractical. This is mainly because establishing predation relationships between the many hundreds of species in an ecosystem requires specialist expertise in species identification and considerable investment in field observation and laboratory experimentation.

The difficulties in deriving ecological networks therefore severely limit our ability to model and predict responses to changes in ecosystem management and any technique which can automate the discovery of plausible trophic links from ecological data is highly desirable.

4. Automated discovery of ecological networks from FSE data and ecological background knowledge

Many forms of Machine Learning, such as Neural Nets and Support Vector Machines, cannot make use of domain knowledge (i.e. ecological knowledge in this study). By contrast, Inductive Logic Programming (ILP) techniques (Muggleton 1991; Muggleton and De Raedt 1994) support the inclusion of such background knowledge and allow the construction of hypotheses that describe structure and relationships between sub-parts. ILP systems use given example observations E and background knowledge B to construct a hypothesis H that explains E relative to B . The components E , B and H are each represented as logic programs. Since logic programs can be used to encode arbitrary computer programs, ILP is arguably the most flexible form of Machine Learning, which has allowed it to be successfully applied in complex problems (Santos et al. 2012; Bohan et al. 2011; Tsunoyama et al. 2008).

In this section we describe an Abductive ILP approach which has been used to automatically generate plausible and testable food-web theories from ecological census data and existing ecological background knowledge. The main role of abductive reasoning in machine learning of scientific theories is to provide hypothetical explanations of empirical observations (Flach and Kakas 2000). Then, based on these explanations, we try to inject back into the scientific theory new information that helps complete the theory. This process of generating abductive explanations and updating theory can be repeated as new observational data become available. The process of abductive learning can be described as follows. Given a theory, T , that describes our incomplete knowledge of the scientific domain and a set of observations, O , we can use abduction to extend the current theory according to the new information

contained in O . The abduction generates hypotheses that entail a set of experimental observations subject to the extended theory being self-consistent. Here, entailment and consistency refer to the corresponding notions in formal logic. Abductive Logic Programming (Kakas et al. 1993) is typically applied to problems that can be separated into two disjoint sets of predicates: the observable predicates and the abducible predicates. In practice, observable predicates describe the empirical observations of the domain that we are trying to model. The abducible predicates describe underlying relations in our model that are not observable directly but can, through the theory T , bring about observable information. Hence, the hypothesis language (i.e. abducibles) can be disjoint from the observation language. We may also have background predicates (prior knowledge), which are auxiliary relations that help us link observable and abducible information.

In many implementations of abductive reasoning, such as that of Progol 5.0 (Muggleton and Bryant 2000), as used in this chapter, the approach taken is to choose the explanation that 'best' generalises under some form of inductive reasoning. This link to induction then strengthens the role of abduction to machine learning and the development of scientific theories. We refer to this approach as Abductive ILP (A/ILP).

A/ILP has been used in a series of studies involving the inference of biological network models from example data. In Tamaddoni-Nezhad et al. (2006) encoding and revising logical models of biochemical networks was done using A/ILP to provide causal explanations of rat liver cell responses to toxins. The observational data consisted of up and down regulation patterns found in high throughput metabonomic data. This approach was further extended by Sternberg et al. (2013), where a mixture of linked metabonomic and gene expression data was used to identify biosynthetic pathways for capsular polysaccharides. In this case ILP was shown to provide a robust strategy to integrate results from different experimental approaches.

A/ILP was also used Tamaddoni-Nezhad et al. (2012) to infer probabilistic ecological networks from the FSE data described in Section 2. The Vortis and Pitfall datasets used for the machine learning were year total data, produced by summing the counts from each sample date, for each taxon in each half-field. This raw data was used to measure a treatment effect ratio: counts from each conventional and GMHT half-field pair were converted into a geometric treatment ratio, as used in Haughton et al. (2003). Counts were log-transformed, using formula $L_{ij} = \log_{10}(C_{ij} + 1)$, where C_{ij} is the count for a species or taxon in treatment i at site j . Sites where $(C_{1j} + C_{2j}) \leq 1$ were removed from the learning dataset (as in Haughton et al. 2003). The treatment ratio, R , was then calculated as $R = 10^d$ where $d = (L_{2j} - L_{1j})$. Following the rationale in Squire et al. (2003), important differences in the count between the two treatments were considered to be greater than 50%. Thus, treatment ratio values of $R < 0.67$ and $R > 1.5$ were regarded as important differences in count with direction of *down* (decreased) and *up* (increased) in the GMHT treatment, respectively. This information on *up* and *down* abundances is considered as our

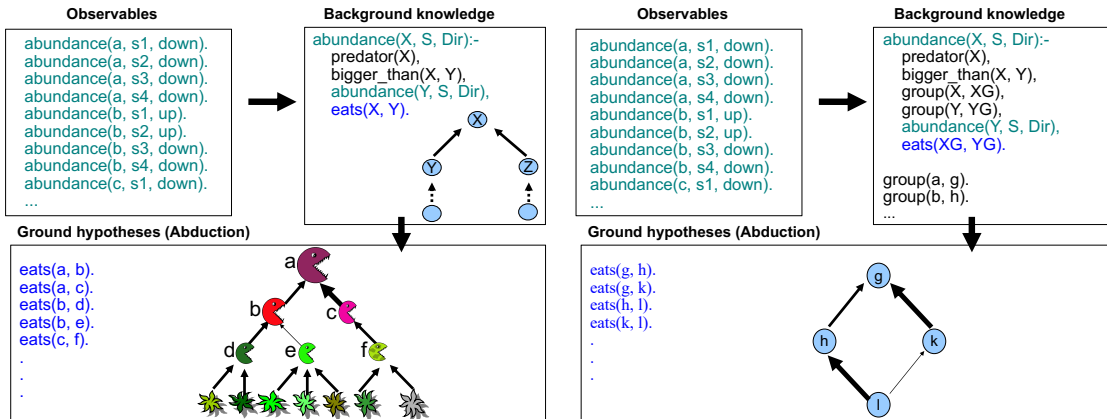


Figure 3 Machine learning of species (left) and functional (right) food-webs from ecological data using Abductive ILP.

observational data for the learning and can be represented by predicate $abundance(X, S, up)$ (or $abundance(X, S, down)$) stating the fact that the abundance of species X at site S is up (or $down$).

The knowledge gap that we initially aimed to fill was a predation relationship between species. Thus, we declare abducible predicate $eats(X, Y)$ capturing the hypothesis that species X eats species Y . It is clear that this problem has properties that require an abductive learning approach such as A/ILP: firstly, the theory describing the problem is incomplete; and secondly, the problem requires learning in the circumstance in which the hypothesis language is disjoint from the observation language.

In order to use abduction, we also need to provide the rules which describe the observable predicate ($abundance$) in terms of the abducible predicate ($eats$):

```

abundance(X, S, Dir):-
    predator(X),
    bigger_than(X, Y),
    abundance(Y, S, Dir),
    eats(X, Y).

```

where Dir can be either up or $down$. This Prolog rule expresses the inference that following a perturbation in the ecosystem (caused by the management), the increased (or decreased) abundance of species X at site S can be explained by X eating species Y , and the abundance of species Y is increased (or decreased). This rule also includes additional conditions to constraint the search for abducible predicate $eats(X, Y)$. These constraints are 1) X should be a predator and 2) X should be bigger than Y . Predicates $predator(X)$ and $bigger_than(X, Y)$ are provided as part of the background knowledge. The 'ecological' background knowledge that a predator should be bigger than a prey was provided by the domain expert.

Given this model and the observable data, the Abductive ILP system Progol 5.0 (Muggleton and Bryant 2000) was used to generate a set of ground abductive

hypotheses in the form of 'eats' relations between species as shown in Fig. 3. These abductive hypotheses are generated by matching observable input against the background knowledge (which includes the rule describing the observable predicate in terms of abducible predicate). In general, many choices for matching could be made, leading to a variety of alternative hypotheses and a preference is imposed by Progol 5 using an information-theoretic criterion known as compression (Muggleton and Bryant, 2000). Here, compression can be defined as $p - n - h$, where p is the number of observations correctly explained by the hypothesis, n is the number incorrectly explained and h is the length of the hypothesis (e.g. 1 for a single fact such as a trophic link). The set of ground hypotheses can be visualised as a network of trophic links (a food-web) as shown in Fig. 4. In this network a ground fact $eats(a, b)$ is represented by a directed trophic link from species b to species a .

A Probabilistic ILP (PILP) approach, called Hypothesis Frequency Estimation (HFE) (Tamaddoni-Nezhad et al., 2012), was used for estimating the probabilities of hypothetical trophic links based on their frequency of occurrence when randomly sampling the hypothesis space. HFE is based on direct sampling from the hypothesis space. In some ILP systems, including Progol 5.0, training examples act as seeds to define the hypothesis space (e.g. a most specific clause is built from the next positive example). Hence, permutation of the training examples leads to sampling from different parts of the hypothesis space. Using this technique, the thickness of trophic links in Fig. 4 (and Fig 5) represent probabilities which are estimated based on the frequency of occurrence from 10 random permutations (a user selected parameter) of the training data (and hence different seeds for defining the hypothesis space).

A probabilistic trophic network can be also represented using standard PILP representations such as SLPs (Muggleton, 1996) or ProbLog (De Raedt et al., 2007). For this we can use relative frequencies in the same way probabilities are used in PILP. We can then use the probabilistic inferences based on these representations to estimate probabilities. For example, the probability $p(abundance(a, s, up))$ can be estimated by relative frequency of hypotheses that imply a at site s is up . Similarly, $p(abundance(a, s, down))$ can be estimated and by comparing these probabilities we can decide to predict whether the abundance is up or $down$.

Species food-web (Fig 4) can be used to explain the structure and dynamics of a particular ecosystem. However, functional food-webs which represent trophic interactions between functional groups of species, might be more important for predicting changes in agroecosystem diversity and productivity (Caron-Lormier et al., 2009). Species in FSE data can be classified into 'trophic-functional types' using general traits that reflect their functional type, primarily resource acquisition and attributes (Caron-Lormier et al., 2009).




Figure 4 Species food-web learned from farm-scale evaluations (FSE) of GMHT crops data collected using Vortis sampling method from 257 fields across the UK. Thickness of trophic links represents probabilities which are estimated using Hypothesis Frequency Estimation (HFE).




Figure 5 Functional food-web learned from FSE data (Vortis). Each group in the functional food-web is represented by a species which can be viewed as an archetype for that functional group.

By assuming that the background knowledge includes information on the functional group of each species, trophic networks for functional groups can be also learned from ecological data using the machine learning approach described above (See Fig. 3). Here we need a rule which describes the observable predicate in terms of eats relation between functional groups:

```

abundance(X,S,Dir):-
    predator(X),
    bigger_than(X, Y),
    group(X, XG),
    group(Y, YG),
    abundance(Y, S, Dir),
    eats(XG, YG).

```

Given this new model and background information, i.e. functional group of species in the form of $group(X, XG)$, trophic networks can be constructed for functional groups in a learning setting similar to the one described above for individual species.

Fig. 5 shows a functional food-web learned from the FSE data (Vortis). This food-web is constructed by learning trophic interactions between functional groups rather than individual species. Each functional group is represented by a species which can be viewed as an archetype for the functional group.

Evaluating food-webs learned from a set of crops on unseen data from a different crop was done by repeatedly constructing food-webs from all crops data, excluding test data from a particular crop and measuring the predictive accuracy on this test data. Fig. 6 show predictive accuracies of Vortis species-based and functional food webs on different crops. The average predictive accuracies (the proportions of correctly predicted left-out test examples) are reported with standard errors associated with each point where 0% to 100% of the training examples are provided.




Figure 6 Predictive accuracies of functional food-web vs. species food-web from cross-validation tests on different crops.

In these experiments, Hypothesis Frequency Estimation (HFE) (Tamaddoni-Nezhad et al., 2012), was used for estimating probabilities of hypothetical trophic links based on the frequency of occurrence from 10 random permutations of the training data. The HFE method was also used in the leave-one-out cross-validation to compare the predictive accuracies of species food-web vs functional food-webs (food-webs shown in Figs. 4 and 5). The experimental materials and methods are described in (Tamaddoni-Nezhad et al., 2013).

According to this figure, the predictive accuracies of the learned food-webs were significantly higher than the default accuracy of the majority class (around 55%). Predictive accuracies for the functional food webs were the same or higher than their species-based counterpart, particularly at low to medium percentages of training examples. This suggests that the functional food webs are at least as accurate as their species-based counterpart, but are much more compact (parsimonious). We also expect the higher predictive accuracy of the functional food-web to be more evident if the food-webs are evaluated on a different agricultural system where different species (not present in the training of species food webs) may exist.

5. Evaluation of the results and subsequent discoveries

The initial species food-webs discovered by machine learning, were examined in Bohan et al. (2011), by domain experts from Rothamsted Research UK and it was found that many of the learned trophic links, in particular those ascribed with high probability by machine learning, are corroborated by the literature. In some cases, novel and high probability links were suggested, and some of these were tested and corroborated by subsequent empirical studies (Davey et al. 2013).

| | | | | | | | | | | | | | | | |
|-----------------------------|---------------------|-------------------|---------------------|-------------------|------------------|---------------|---------------|-------------|----------------------|------------------------|----------------|---------------|------------|------------------------|--------------|
| Agonum dorsale | Anthracoris nemorum | Bembidion lampros | Bembidion lunulatum | Bembidion obtusum | Cimicidae nymphs | Curculionidae | Entomobryidae | Isotomidae | Lepthyphantes tenuis | Liccoris tripustulatus | Miridae nymphs | Ortus vicinus | Poduridae | Scotopostellus affinis | Sminthuridae |
| | | | | | | | 9 [13] | 4 | 5 [6] | | | | | | 10 [13] |
| Bembidion aeneum | | | | | | | 10 [11] | | 9 [6] | | | | 9 [11] | | |
| Bembidion biguttatum | | | | | | | | | | | | | | | 10 [11] |
| Bembidion guttula | | | | | | | 7 [11] | 10 [11] | 9 [6] | | | | | | |
| Bembidion lampros | | | | | | | 10 [11] | 10 [11] | 10 [6] | | | | 10 [11] | 9 | 10 [11] |
| Bembidion obtusum | | | | | | | 10 [11] | 10 [11] | 10 [6] | | | 10 | | | 10 [11] |
| Bembidion quadrimaculatum | | | | | | | 10 [11] | 10 [11] | 9 [6] | | | 10 | | | 9 [11] |
| Brachycellus verbasci | | | | | | | | | | | | | | | |
| Carabid larvae | | 9 | 10 | 3 | | | 10 [2] | 10 [2] | 9 | | 10 | | 9 [2] | | 10 [2] |
| Clivina fossor | | | | | | | 7 [12,5] | | 2 | | | | | | |
| Coccinellid larvae | | 9 | | | | | 10 [14,16] | 9 [14,16] | | | 1 | | 10 [14,16] | | 9 [14,16] |
| Coccinella septempunctata | | | | | | | | 10 [14,16] | | | | | | | |
| Diomus linearis | | | | | | | 10 [11] | 7 | | | | | 3 | | |
| Loricera pilicornis | | | | | | | | | | 9 | | | | | |
| Nabidae nymphs | | | | 3 | | | | 10 | 7 | | 2 [10,7] | | | | 10 [9] |
| Nebria brevicollis | | 10 | | | | | | 10 [15] | | | | | | | |
| Notophilus biguttatus | | | | | | | 10 [4,11,3] | 10 [4,11,3] | | | | | | | 10 [4,11,3] |
| Proxya quatuordecimpunctata | | | | | | | 10 [14,16] | | 10 [16] | | | | | | |
| Saldula saltator | | | | | | | 10 [8] | | | | | | | | |
| Trechus quadristriatus | | 9 | | | | | 9 [15,3] | 4 [15,3] | 9 [6] | | 2 | | | | 10 [15,3] |
| Trechus secalis | | | | | | | | 8 [15,3] | | | | | | | |

Figure 7 Manual corroboration of trophic links for some prey (columns) and predator (rows) species combination from Fig. 4. Each pairwise hypothesised link has a strength (i.e. frequency between 1 to 10) followed by references (in square brackets) in the literature (see Appendix 1 of (Tamaddoni-Nezhad et al., 2012)) supporting the link. Multiple references are indicated by yellow and green ellipses and potential novel hypotheses by dashed red ellipses.

Manual examination of the food-webs was used to corroborate some of known trophic links and also to identify potential novel hypotheses as shown in Fig. 7. However, manual corroboration of hypothetical trophic links is difficult and requires significant amounts of time and is error prone. Hence, a text-mining technique was adopted (Tamaddoni-Nezhad et al., 2013) for automatic corroboration of hypothetical trophic links from ecological publications. This was particularly useful for larger food-webs from merged Vortis and pitfall data.

Fig. 8 illustrates how a literature network can be generated based on the co-occurrences of predators/prey species in the relevant context, directly from the literature. The pairs of species (from a given food web) and the interaction lexicons (from a dictionary file) are used to generate queries. Then the text-mining module searches through the text of available publications to match each query. The publications can be in a local database or accessed via a search engine (e.g. Google Scholar). The output of the text-mining for each query is the number of publications that matched that query (number of hits). The output for a whole food web can be represented by a literature network in which the number associated with each edge is related to the number of papers where the co-occurrences of the predator / prey species have been found with at least one trophic interaction lexicon (eat, feed, prey, or consume). We have shown that the frequencies of trophic links (using HFE) are significantly correlated with the total number of hits for these links in the literature networks (Tamaddoni-Nezhad et al., 2013). Moreover, the proposed approach was used to successfully identify hypothetical trophic relations for which there are little or no information in the literature (potential novel hypotheses).

The manual corroboration table (Fig. 7) represents prey (columns) and predator (rows) species combination from the Vortis food-web. Each pairwise hypothesised link has a strength (i.e. frequency between 1 to 10, from the HFE method) followed by references (in square brackets) in the literature (see Appendix 1 of (Tamaddoni-Nezhad et al., 2012)) supporting the link. This table shows that many of the links, suggested by the model, are corroborated by the literature. In particular, links in the model ascribed with high frequency correspond well with those having multiple references in the literature. For example, there are 15 links with more than two references and 8 of these are with frequency 10 and from these all the three links with three references (marked by green circles) have frequency 10. In addition, there are also highly frequent links with no references in the literature, and these could potentially be novel hypotheses for future testing with targeted empirical data. For example, one surprising result, was the importance of carabid larvae as predators of a variety of prey and in some cases with no reference in the literature (see Fig 7). As another example, some species of spiders appeared as prey for other predators); a result that was unexpected because spiders are obligate predators. This hypothesis was tested in a subsequent study using molecular analysis of predator gut contents and it was found that this hypothesised position in an animal-animal network is correct (Davey et al. 2013), and spiders do appear to play an important role as prey at least for part of the agricultural season. Thus, even though some of the hypothesised links were unexpected, these were in fact confirmed later and this provided an extremely stringent test for this human- machine scientific discovery approach.




Figure 9 Pairwise comparisons of structural properties of individual crop food-webs between conventional and GMHT managements (**a,b**, beet; **c,d**, maize; **e,f**, spring oilseed rape; **g,h**, winter oilseed rape). C represents network connectance; ϕ , core link density; R_R , robustness via random removal; R_T , robustness via targeted removal of highest degree nodes as described in Ma et. al. (2019).

The food-webs constructed and validated using this human-machine discovery approach were also the basis of a recent study revealing resilience of agro-ecosystems to changes in farming management using GMHT crops. Ma et. al. (2019) constructed replicated food webs using the merged Vortis and pitfall food-web, populated on the basis of the sampled taxonomic and abundance information of each half of the split-field in FSE and obtained a total of 502 food webs (251 conventional and 251 GMHT). A network analysis approach was used to characterize the structural properties of all the individual food-webs. The network analysis metrics include: C , connectance; ϕ , core link density; Core size; R_R , robustness via random removal; R_T , robustness via targeted removal of highest degree nodes, as defined in Ma et. al. (2019). Each metric is averaged across all webs of a given variety and normalized by its overall range. The effects of crop type can be visualized by comparing results from conventional crops horizontally as shown in Fig. 9. As shown in this figure, food web properties varied significantly between crop types. However, this figure suggest that the food web properties remain unaltered between conventional and GMHT food webs. The network analysis approach by Ma et. al. (2019) also revealed that network-level responses of GMHT crops are remarkably similar in their composition and responses to simulated trajectories of species removals, to their conventional counterparts. These results suggest that that crop type was by far the dominant driver of differences in web structure and robustness, across several organizational levels, ranging from sub-structural to whole-network attributes; inter-annual variation is probably greater than differences between conventional and GMHT.

6. Conclusions

In this chapter we have demonstrated how a combination of comprehensible machine learning, text-mining and expert knowledge was used to automatically generate plausible and testable food-web hypotheses from ecological census data. The logic-based machine learning included human-encoded ecological background knowledge, e.g. size relationship between predator and prey and taxonomical functional types. Text-mining from scientific publications was initially used to verify machine-learned hypotheses, but it was also useful for identifying potential novel hypotheses, i.e. high probability hypotheses suggested by machine learning with no references in the literature. The results included novel food-web hypotheses, some confirmed by subsequent experimental studies (e.g. DNA analysis of gut contents) and published in scientific journals. This case study shows the potentials of human-machine collaboration / communication for the purpose of hypothesis generation in scientific discovery.

Fig. 10 shows the cycle of hypothesis generation and experimentation in (biological) scientific discovery. In this cycle, machine learning is usually used

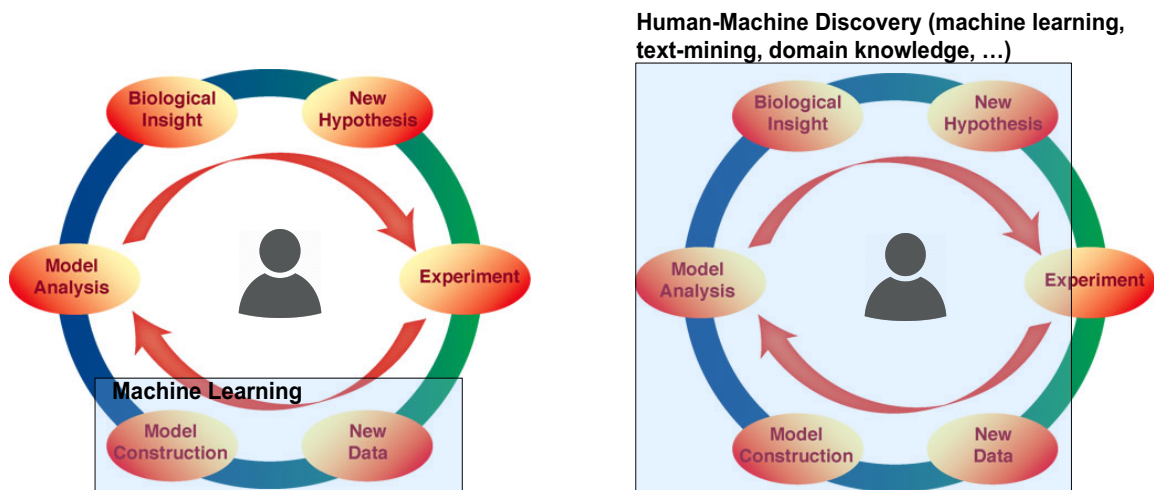


Figure 10 Machine Learning vs Human-Machine Discovery

for 'Model Construction' from 'New Data'. However, the purpose of Human-Machine discovery is to also automate other steps of this cycle by combining machine learning, text-mining and domain knowledge, as in the case study described in this chapter.

We argue that with ever-growing amount of human knowledge and empirical data as well as advances in AI, human-machine discovery, where humans and computers jointly develop and evaluate scientific theories, will be important for the advancement of the science in future.

References

- Bohan, D. A., Boffey, C. W. H., Brooks, D. R., Clark, S. J. , Dewar, A. M., Firbank, L. G., Haughton, A. J., Hawes, C., Heard M. S. and May, M. J. (2005) Effects on weed and invertebrate abundance and diversity of herbicide management in genetically modified herbicide-tolerant winter-sown oilseed rape. *Proceedings of the Royal Society B: Biological Sciences*, 272, 463-474.
- Bohan, D. A., Caron-Lormier, G., Muggleton, S., Raybould, A., and Tamaddoni-Nezhad, A. (2011) Automated Discovery of Food Webs from Ecological Data Using Logic-Based Machine Learning, *PloS One*, 6, e29028.
- Caron-Lormier G, Bohan D.A, Hawes C, Raybould A, Haughton AJ et al. (2009) How might we model an ecosystem? *Ecol Model* 220: 1935-1949.
- Cohen, J.E., Schittler, D.N., Raffaelli, D.G., and Reuman, D.C. (2009) Food webs are more than the sum of their tritrophic parts. *Proc. Natl. Acad. Sci. USA* 106, 22335–22340.
- Champion, G. T., May, M. J., Bennett, S., Brooks, D. R., Clark, S. J., Daniels, R. E., Firbank, L. G., Haughton, A. J., Hawes, C., Heard, M. S., Perry, J. N., Randle, Z., Rossall, M. J., Rothery, P., Skellern, M. P., Scott, R. J., Squire,

G. R. and Thomas, M. R. (2003) Crop management and agronomic context of the Farm Scale Evaluations of genetically modified herbicide-tolerant crops. *Philosophical Transactions of The Royal Society of London, Series B*, 358, 1801–1818.

Davey, J. S., Vaughan, I. P., Andrew King, R., Bell, J. R., Bohan, D. A., Bruford, M. W., Holland, J.M. and Symondson, W.O. (2013). Intraguild predation in winter wheat: prey choice by a common epigeal carabid consuming spiders. *Journal of Applied Ecology*, 50(1), 271-279.

De Raedt, L., Kimmig, A., Toivonen, H. (2007) Problog: A probabilistic prolog and its applications in link discovery. In: Lopez de Mantaras, R., Veloso, M.M. (eds.) *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, pp. 804–809

Dickinson G, Murphy K (1998) *Ecosystems: A Functional Approach*. London: Routledge. 190 p.

Flach, P., & Kakas, A.C. (2000). Abductive and inductive reasoning: Background and issues. In: Flach, P.A., & Kakas, A.C., editors, *Abductive and Inductive Reasoning, Pure and Applied Logic*. Kluwer.

Haughton, A., Champion, G., Hawes, C., Heard, M., Brooks, D., Bohan, D., Clark, S., Dewar, A., Firbank, L. and Osborne, J. (2003) Invertebrate responses to the management of genetically modified herbicide-tolerant and conventional spring crops. II. Within-field epigeal and aerial arthropods, *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 358, 1863-1877.

Ings, T.C., Montoya, J.M., Bascompte, J., Bluthgen, N., Brown, L., Dormann, C.F., Edwards, F., Figueroa, D., Jacob, U., Jones, J.I., Lauridsen, R.B., Ledger, M.E., et al. (2009) Ecological networks - beyond food webs. *J. Anim. Ecol.* 78, 253-269.

Kakas, A.C., Kowalski, R.A. and Toni, F. (1993) Abductive Logic Programming. *Journal of Logic and Computation*, 2(6), 719–770.

Kakas, A.C. and Riguzzi, F. (2000). Abductive concept learning. *New Generation Computing*, 18, 243–294.

Ma, A., Lu, X., Gray, C., Raybould, A., Tamaddoni-Nezhad, A., Woodward, G., & Bohan, D. A. (2019). Ecological networks reveal resilience of agroecosystems to changes in farming management. *Nature Ecology & Evolution*, 3(2), 260-264.

Michie, D. (1988). Machine learning in the next five years. In *Proceedings of the third European working session on learning*, pp. 107–122. Pitman.

Muggleton, S.H. (1991) *Inductive Logic Programming*. *New Generation Computing*, 8(4):295–318.

- Muggleton SH (1996) Stochastic Logic Programs. In: de Raedt L, ed. *Advances in Inductive Logic Programming* IOS Press. pp 254–264.
- Muggleton, S. and Bryant, C. (2000) Theory completion using inverse entailment. In *Proc. of the 10th International Conference on Inductive Logic Programming*, 130-146.
- Muggleton, S.H. and De Raedt, L. (1994) Inductive logic programming: Theory and methods. *Journal of Logic Programming*, 19,20:629–679.
- Odum E (1971) *Fundamentals of Ecology*, 3rd Edition. New York: Saunders Press.
- Perry, J. N., Rothery, P., Clark, S. J., Heard, M. S. and Hawes, C. (2003) Design, analysis and statistical power of the Farm Scale Evaluations of genetically modified herbicide-tolerant crops. *J. Appl. Ecol.* 40, 17–31.
- Pocock, M.J.O., Evans, D.M. and Memmott, J. (2012) The Robustness and Restoration of a Network of Ecological Networks. *Science*, 335, 973–977.
- Reuman, D.C., Mulder, C., Banašek-Richter, C., Cattin Blandenier, M.-F., Breure, A.M., Den Hollander, H.A., Kneitel, J.M., Raffaelli, D., Woodward, G. and Cohen, J.E. (2009) Allometry of body size and abundance in 166 food webs. *Adv. Ecol. Res.* 41, 1-44.
- Santos, J.C.A., Nassif, H. and Muggleton, S.H., Sternberg, M.J.E., Page, C.D. (2012) Automated identification of protein-ligand interaction features using inductive logic programming: A hexose binding case study. *BMC Bioinformatics*, 13(162).
- Squire GR, Brooks DR, Bohan DA, Champion GT, Daniels RE et al. (2003). On the rationale and interpretation of the farm-scale evaluations of genetically-modified herbicide-tolerant crops. *Phil Trans Roy Soc B* 358: 1779-1800.
- Sternberg, M.J.E., Tamaddoni-Nezhad, A., Lesk, V.I. , Kay, E., Hitchen, P.G., Cootes, A., Alphen, L.B., Lamoureux, M.P., Jarrell, H.C., Rawlings, C.J., Soo, E.C., Szymanski, C.M., Dell, A., Wren, B.W. and Muggleton, S.H. (2013) Gene function hypotheses for the campylobacter jejuni glycome generated by a logic-based approach. *Journal of Molecular Biology*, 425(1):186–197.
- Stokes, J. M., Yang, K., Swanson, K., Jin, W., Cubillos-Ruiz, A., Donghia, N. M., MacNair, C.R., French, S., Carfrae, L.A., Bloom-Ackerman, Z. and Tran, V.M. (2020). A deep learning approach to antibiotic discovery. *Cell*, 180(4), 688-702.
- Tamaddoni-Nezhad, A., Bohan, D., Raybould, A. and Muggleton, S. (2012) Machine learning a probabilistic network of ecological interactions, In

Proceedings of the 21st International Conference on Inductive Logic Programming, LNAI 7207, 332-346.

Tamaddoni-Nezhad, A., Chaleil, R., Kakas, A. and Muggleton, S.H. (2006) Application of abductive ILP to learning metabolic network inhibition from temporal data. *Machine Learning*, 64:209–230.

Tamaddoni-Nezhad, A., Chaleil, R., Kakas, A., Sternberg, M., Nicholson, J., and Muggleton, S.H. (2007) Modeling the effects of toxins in metabolic networks. *IEEE Engineering in Medicine and Biology*, 26:37–46.

Tamaddoni-Nezhad, A., Milani, G., Raybould, A., Muggleton, S., Bohan, D. (2013) Construction and validation of food-webs using logic-based machine learning and text-mining. *Advances in Ecological Research* 49, 225–289

Tsunoyama, K., Amini, A., Sternberg, M.J.E. and Muggleton, S.H. (2008) Scaffold hopping in drug discovery using inductive logic programming. *Journal of Chemical Information and Modelling*, 48(5):949–957.

Woodward, G., Brown, L. E., Edwards, F. K., Hudson, L. N., Milner, A. M., Reuman, D. C., & Ledger, M. E. (2012). Climate change impacts in multispecies systems: drought alters food web size structure in a field experiment. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367(1605), 2990-2997.