Machine Education: Designing semantically ordered and ontologically guided modular neural networks

Hussein A. Abbass, Sondoss Elsawah *School of Engineering & IT*
University of New South Wales
Canberra, Australia
{s.elsawah,h.abbass}@adfa.edu.au Eleni Petraki *Faculty of Education*
University of Canberra
Canberra, Australia
Eleni.Petraki@Canberra.edu.au Robert Hunjet *Land Division*
Defence Science and Technology
Adelaide, Australia
robert.hunjet@dst.defence.gov.au

Abstract—The literature on machine teaching, machine education, and curriculum design for machines is in its infancy with sparse papers on the topic primarily focusing on data and model engineering factors to improve machine learning. In this paper, we first discuss selected attempts to date on machine teaching and education. We then bring theories and methodologies together from human education to structure and mathematically define the core problems in lesson design for machine education and the modelling approaches required to support the steps for machine education. Last, but not least, we offer an ontology-based methodology to guide the development of lesson plans to produce transparent and explainable modular learning machines, including neural networks.

Index Terms—Curriculum Design, Machine Teaching, Machine Education, Transparent and Explainable Modular Neural Networks

I. Introduction

WITH recent advances in deep learning and a new-found abundance of data on massive scales, researchers have started to feel the need to follow structured methodological approaches to teaching machines in similar ways humans follow structured teaching approaches when teaching other humans, leading to the fields of machine teaching [1], curriculum design for machines [2], and machine education [3].

A common aspect of this domain's literature is that it is driven by machine learners alone with a focus on data engineering and model engineering to improve machine learning. As such, except for a couple of recent attempts [4], [5], [6], the literature on machine teaching and education has not taken inspiration from the more mature literature on human teaching and education which houses a wealth of pedagogical frameworks, philosophies and theories that could benefit the machine education field.

In this paper, we will present an initial approach to connect threads from the human education literature to machine education. The remainder of this paper is structured as follows. Section II acts as a short background section on current work on machine teaching and education. Section III presents key concepts in pedagogy and curriculum design methodologies in human education, followed by the first attempt to structure machine education in a manner that leverages the significant literature on human education, thereby addressing significant methodological gaps in the machine education literature. We then conclude the paper with a discussion on open challenges in the field.

II. Machine Teaching and Education

The emerging literature of machine-teaching can be traced back to Elman's work in 1993 [7], which promoted the need to 'start small ' and structure learning experiences from simple to more complex. Elman systematically increased the memory of the learner while fixing the complexity of the input to the learner, and was able to teach the machine to learn a 'semi-realistic artificial language' with increasing levels of competencies. These ideas have anchored the roots of machine-teaching in the literature, which started to emerge as a field only in the last few years. However, similar to Elman, the work is mostly driven by experts in machine learning with intimate knowledge of algorithms, simple tasks, and, in most cases, exiguous grounds in teaching methodologies. Moreover, in the field of evolutionary-learning, other than limited recent attempts [5], the research area is in its infancy. To the best of our knowledge, attempts in the evolutionary neural network literature are rare.

Clayton and Abbass [5] designed a curriculum suitable for reinforcement learning used within an autonomous agent in shepherding tasks. They extended Dick and Carey's Model for Systematic Instructional Design (SID) [8], [9], which offers the following ten steps: identification of an instructional goal, conducting instructional analysis, identification of entry behaviours and characteristics, write-up of performance objectives, development of criterion-referenced assessments, development of instructional strategies, development of and/or select instruction, development of and carrying out of formative evaluations, revising the instruction, and conducting summative evaluation.
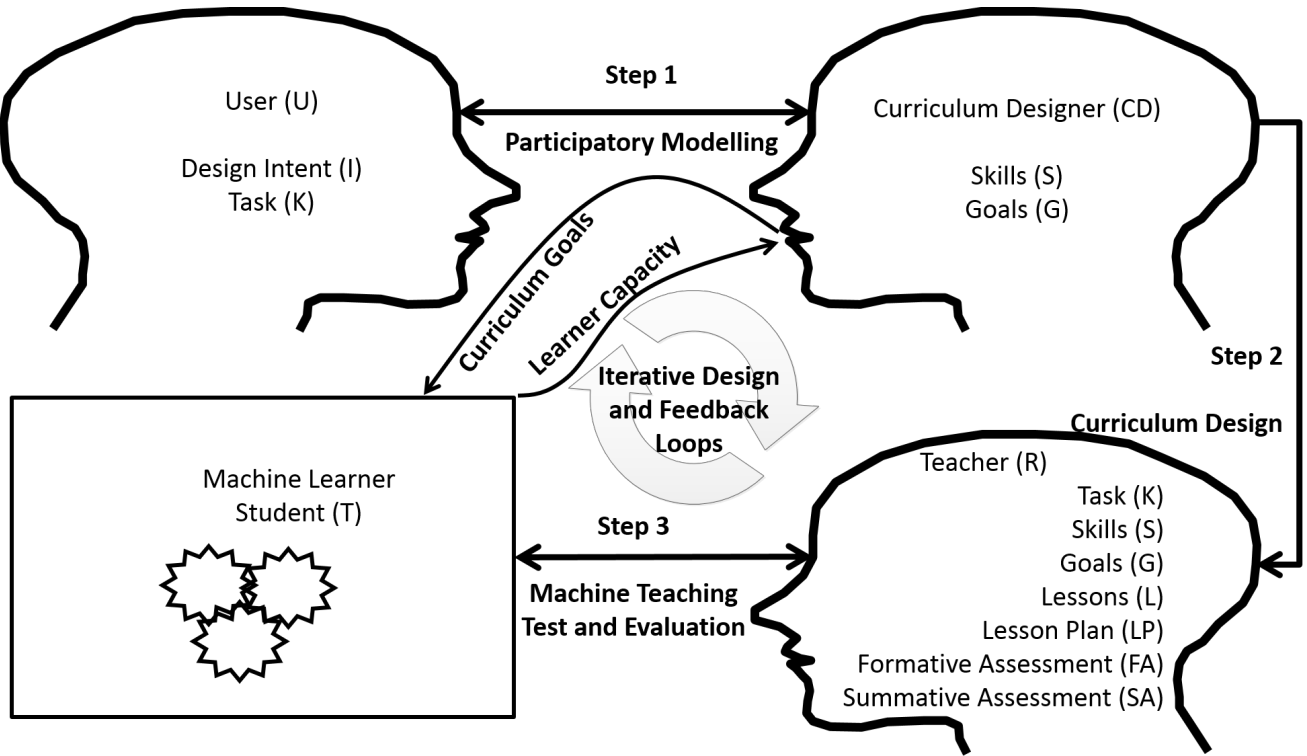
Fig. 1. The methodological framework for machine education.

Gee and Abbass [6] offered a complete methodology to teach machines by extending George's classical curriculum design [10] with task analysis to suit machine education. They used a supervised learning approach to the shepherding problem to demonstrate the methodology and were successful in learning primitive shepherding behaviours.

Zhu et al. [4] presented a preliminary high-level study to ground machine teaching in some concepts in human education. They introduced eight dimensions to characterize the different threads of research that have either taken place or could take place in the area of machine teaching: Human vs. machine; Teaching signal; Batch vs. sequential teaching; Model-based vs. model-free; Student awareness; One vs. many; Angelic vs. adversarial; and Theoretical vs. empirical.

Aside from the above, all other research on machine teaching in the literature relies on classic machine learning and deep learning, merely focusing on task decomposition and data engineering.

III. MACHINE TEACHING AND EDUCATION

The primary challenge in transforming concepts from the human education literature to machine education is the differences between current state-of-the-art in machine cognition and human cognition. Humans can understand and act on ambiguous language sitting at a variety of levels of abstraction. Human teachers follow pedagogical theories that allow them a great deal of freedom in operationalizing the concepts and the fluidity to adopt these concepts differently based on a learner's needs. Machines require a more formal representation of these concepts in the form of mathematical models. To ensure the trustworthiness of the machine learners produced

through machine education, assessments, for example, need to be captured in formal models that are mathematically and logically verifiable. At this stage, we need to transform the above concepts mathematically adjust them for a machine context. We will start with the following problem definition, where we will distinguish between machine education, which is concerned with the design and implementation of the process, and machine teaching, which is concerned with the delivery of materials to the machine.

Given a user, U, wishing to design an Artificial Intelligence (AI) student, T, to fulfil U's design intent, I, by performing task, K, a curriculum designer, CD, aims to transform I into a set of concrete learning skills, S, and a set of goals, G, that concretely measure that T developed S successfully. Given S and G, CD needs then to define the set of lessons, L, lesson plan, PL, a set of formative assessments, FA, and a set of summative assessments, SA, to teach S to T such that all goals in G are achieved as measured by the formative and summative assessments, FA and SA, respectively. If a syllabus, Y, is defined by $Y = (L, LP, FA, SA)$, a learner-centric curriculum is defined by the tuple $C = (T, K, S, G, Y)$. Machine education is to design C and select the teacher, R, while machine teaching is to select and implement Y such that, R delivers Y to T, T develops skills S to perform K, so that the goals G are met as measured by FA and SA.

We define three broad steps representing the three challenges that need to be overcome to move from the need analysis for curriculum design all the way to a functional syllabus that could be used to teach machine learners. These three steps are presented in the conceptual diagram in Figure 1 and are discussed below:

1) Our first step in designing theoretical constructs for machine teaching is to identify an appropriate modelling paradigm to operationalize the human education concepts presented in the previous section and transform them into objective and quantitative models. Formally, this step could be represented as follows:

$$Step1 : (I \wedge K) \rightarrow (S \wedge G)$$

This step will be addressed in the next subsection, which will introduce participatory modelling and the use of system dynamics as a modelling framework suited for representing the complexity that could arise from complex problems in machine learning. The primary aim of these participatory modelling approaches is to transform user requirements in terms of what the user expects the machine learner to learn (need–analysis in curriculum design) into formal specifications of the learning outcomes, formulation of objectives, and models for testing and evaluation of the educational journey of a machine learner.

2) The second step is the design of syllabus, and in particular the decomposition of the overall curriculum into lessons. The following sub-section will focus more on the syllabus and will present the concept of scaffolding as a means of decomposing a learning task into chunks. We will then require that each chunk is associated with appropriate semantics, such that the orders of chunks follow appropriate semantic orders that allow the overall curriculum to be semantically ordered and fit within an appropriate ontological framework.

$$Step2 : (S \wedge G) \rightarrow Y$$

3) The third step focuses on the curriculum design itself. This step is not discussed in this paper is and the context of our future work.

$$Step3 : (S \wedge G \wedge Y) \rightarrow C$$

A. Participatory and system dynamics modelling

The aim of this step is to transform a user intent and tasks into precise skills and goals; that is, $(I \bigwedge K) \rightarrow (S \bigwedge G)$. Participatory modelling is a purposeful learning process for action that engages the implicit and explicit knowledge of problem actors to create formal and shared representations of the problem of interest [11]. The key distinction of participatory modelings from other participatory and actionable science research approaches is its unique focus on the utilization of different types of models (i.e. conceptual and quantitative) and modelling approaches (e.g. system dynamics) to structure the learning process, organize and integrate knowledge (i.e. sources and forms), and package the produced knowledge into useful artefacts. This modelling focus offers a clear stepwise methodological framework to capture an actor's mental model (i.e. perceptions, knowledge, requirements) and transform them into operational outcomes [12].

There exist several typologies for characterising the expected outcomes of employing participatory modelling [13]. Broadly speaking, operational outcomes can be categorised into: cognitive (e.g. increase in the level of actor's knowledge or skills), social (e.g. establishing formal rules to organize actors' interactions), and normative outcomes (e.g. development of methods for decision analysis). Participatory modelling has been applied in a variety of domains achieved a wide range of purposes, including decision making [14], deliberation and negotiation [15], futuristic system design [16], and to much less extent support formal education and curriculum design. Various participatory modelling frameworks exist that propose specific process steps and combinations of qualitative and quantitative modelling methods [12]. Major differences among existing participatory modelling frameworks relate to the modelling methods included, such as agent-based modelling (see for example [17]) or system dynamics modelling (see for example [18]. Consideration of the context in the design of participatory modelling processes is a critical factor [18].

Regardless of the detailed differences stemming from the choice of the modelling framework, a typical participatory modelling approach proceeds through iterations of the following mechanisms [19]: problem scoping and context setting, process design, model development, and feedback, monitoring and evaluation.

Based on systems thinking and control theory, system dynamics encompasses a range of methods including: conceptual and graphical models (e.g. influence diagrams, stock and flow diagrams), quantitative models (e.g. simulation, games [20]), and participatory methods (e.g. groups model building [21]). At the core of the conceptual basis of system dynamics, learning plays a central role in shaping the development of system dynamics methods and praxis. Learning, about complex problems and the mental models that individuals and groups have about these problems, is recognised as a key rationale and outcome for using system dynamics methods [22]. A typical system dynamics process progresses through the following steps [23]: (1) problem scoping, (2) conceptual modelling, (3) model formulation and implementation, (4) model use.

The process of operationalizing human education concepts into machine education is mainly concerned with addressing the fundamental question of 'how can we teach machines effectively'? In other words, what is the relative effectiveness of different pedagogical theories, strategies, tactics, syllabus, and the overall curriculum design in producing the desirable learning outcomes? This is a question that can be addressed through theoretical arguments, but equally through experiments based on simulation of system dynamics models. Conceptual system dynamics techniques (such as causal loop diagrams and stocks and flows diagrams) capture the implicit and explicit assumptions about the learning process, with particular focus on interdependency and feedback interactions that shape the learning process and determine its outcomes. Data to support the development and validation of these assumptions, referred to as dynamic hypotheses, are gathered through multiple sources, such as expert opinions, literature surveys, and empirical data [23]. The transformation of the conceptual model into a quantitative system dynamics model allows for testing the

effectiveness of different learning theories and implementation scenarios. System dynamics takes the view that any type of dynamic abilities (e.g. skills, knowledge) can be represented as a stock (a rectangle), or a set of related stocks (See Figure 2 for an example), that accumulates or depletes over time in effect of changes in inflow (e.g. knowledge acquisition, skill building) and outflow (e.g. unlearning skills).

The figure depicts a simplified model of the curriculum process, where a gap between the current competency of the machine and its target skill level has been identified. The speed of the learner along with learning outcomes drawn from a lesson pool form the basis upon which skills are built. As skills continue to flow in the mental model of the learner, the first level of skill starts to increase. When the target level of skill is reached, the learner is ready and the skills start to build up in the next level. The control loop continues through a set of adjusting actions that control the draw of lessons from the lesson pool based on observations arising from how the learner is meeting a target skill level. This stock-and-flow diagram has multiple uses. First, it forms the basis for the controller of the machine teaching process that decides when and how much certain data flow to the machine for training. Second, it is a diagnostic tool that helps to identify when the targets for certain skills are not met. Third, it could communicate to, and guide, the teacher (human or machine) in the machine teaching process. The next sub-section will zoom in on stock rectangle labelled as "lessons pool" to look into how the overall syllabus gets decomposed into lessons.

B. Scaffolding and Syllabus Design

The aim of this step is to transform the skills and goals from the previous steps into a syllabus; that is, $(S \wedge G) \rightarrow Y$. Recalling that a syllabus is made of lessons, lessons plan, formative, and summative assessment; that is, $Y = (L, LP, FA, SA)$. A key challenge in this step is to transform the skills and goals into lessons and decomposing these lessons in a way that the machine could be taught.

Machine teaching attempts to learn the data set that generated a specific model. This process is a form of inverting a model by projecting it back to the feature space and the data sample that was used to parameterise the model. Equally, if we have a learning outcome that needs to be achieved, machine teaching will generate the learning lessons necessary to reach this learning outcome. For artificially supervised learners, this step is about the generation of the experience, observations, or training samples that are needed to lead to a specific model.

In this section, we propose Incremental Semantic Machine Teaching (ISMT), where the objective is to transform the discovery of the sample to a learning process, whereby the learner starts with a random or targeted sample (called a chunk) then learns to add chunks as necessary to move the learner as close as possible to achieving the overall learning outcomes. In this incremental process, the choice of the sample is guided by the semantics to be encoded within intermediate models.

The concept of scaffolding [24] in human education involves the decomposition of the curriculum into chunks and the sequencing of these chunks in the learning process to incrementally evolve the learner's knowledge. Hammond and Gibbons [24] proposed scaffolding at the macro "designed-in" level. They decomposed the curriculum design process into seven stages: determining student's experience and prior learning, selection of tasks, sequencing of tasks, participant structures, semiotic systems, mediational texts, and metalinguistic and metacognition awareness. We will focus on the formalism for the first four stages and postpone the latter stages to our future work.

Define a target function $\vec{y} = f(\vec{w}, \vec{x})$, where $\vec{w}$ are the weights parameterizing the function. The Machine Learning (ML) problem can be defined as:

*Definition 1 (***ML Problem***):* Given a training set $\Gamma = (\vec{x}, \vec{y})$, find $\vec{w}$ such that $\hat{\vec{y}} = f(\vec{w} \otimes \vec{x})$ and the loss function, $Err(\hat{\vec{y}}, \vec{y})$, is minimum.

In the above definition, $\vec{x}$ represents the inputs (an $m \times n$ matrix of m observations and n, preferably independent, variables) and $\vec{y}$ represents the true output vector. Together, they form the sample, Γ, used for building model $\hat{\vec{y}}$. The error function $Err(\hat{\vec{y}}, \vec{y})$ can be arbitrarily chosen to be any type of loss function.

The Machine Teaching (MT) problem is the dual of the ML problem above, where the problem is to find an estimate of the sample $\hat{\Gamma}$, that if it was used in place of Γ, it would generate an equivalent model $\hat{\vec{y}}$. The problem can be defined formally as follows:

*Definition 2 (***MT Problem***):* Given a model $\hat{\vec{y}} = f(\vec{w} \otimes \vec{x})$ and its parameters $\vec{w}$, find $\hat{\Gamma}$ such that the loss function, $Err(\hat{\vec{y}}, \vec{y})$, is minimum.

One of the greatest statisticians of his time wrote: "..., *the object of statistical methods is the reduction of data. A quantity of data, by which usually by its mere bulk is incapable of entering the mind, is to be replaced by relatively few quantities which shall adequately represent the whole, of which, in other words, shall contain as much as possible, ideally the whole, of the relevant information contained in the original data*" (p278, [25]).

Even though ML attempts to find this reduced form by describing a model of the data, so often some form of parsimonious pressures are needed to further ensure that the model is the shortest description possible [26]. Parsimonious pressures in fields such as neural networks are achieved through a regularization term that gets added to the loss function.

Similar to ML where the parsimonious pressure is to find the shortest model, the parsimonious pressures leading to efficiencies in MT (we will call it EMT for efficient MT) is to find the smallest dataset. This may not always be preferable. For example, the most comprehensible dataset to teach a human may not be practically the most efficient. Redundancy and repetitions are sometimes useful tools to reinforce the message for the learner.

Nevertheless, to teach a machine, EMT contributes to efficiency of resources and ensures the machine is taught in the most effective manner. An advantage of EMT is that it reduces demands on the machine's working memory. The definition of EMT is:
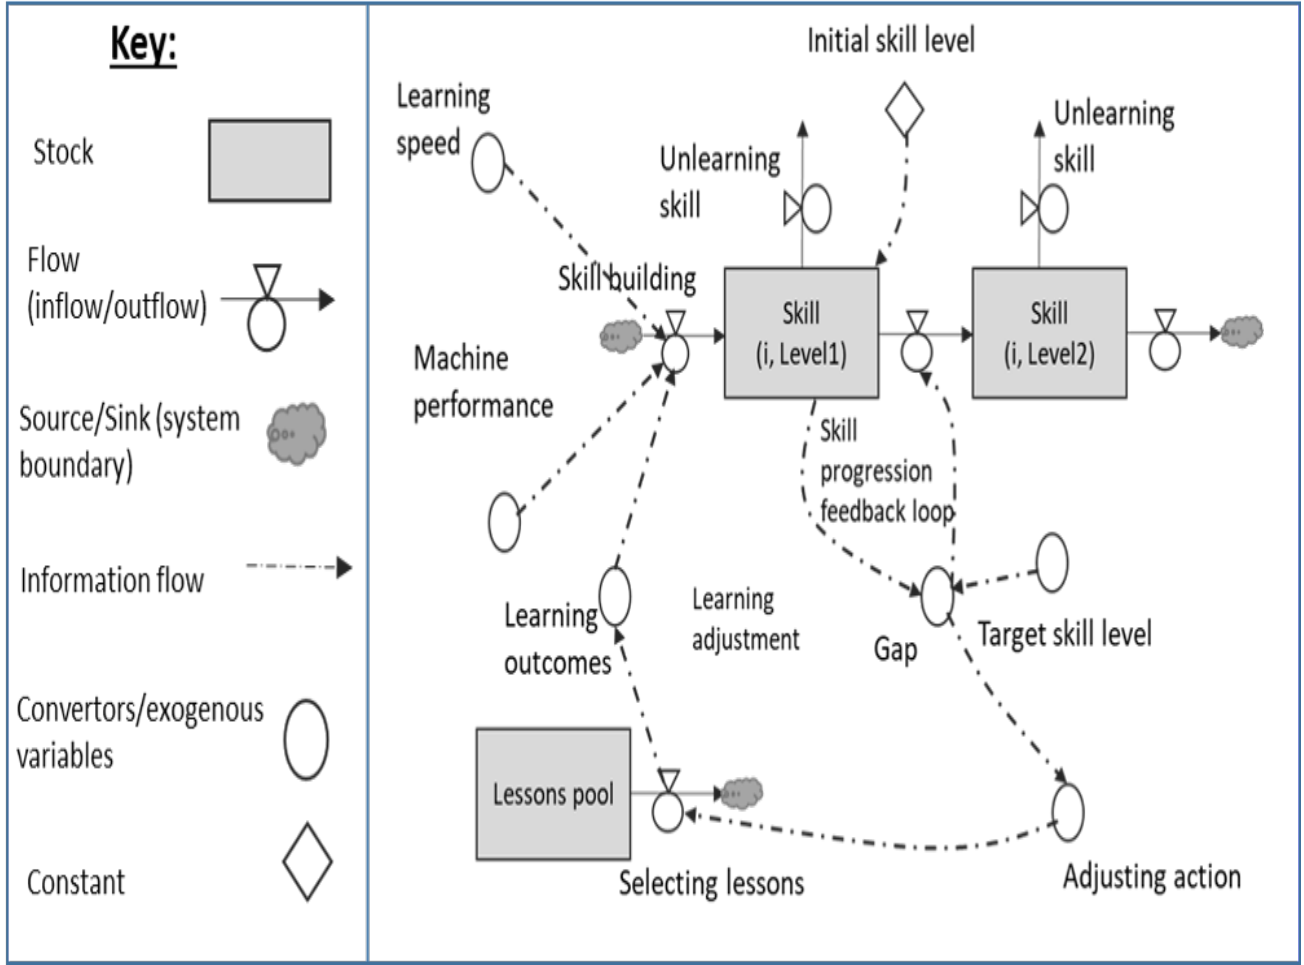
Fig. 2. High level stock and flow diagram of the progression between skill levels.

Definition 3: [**EMT Problem**] **Given a model** $\hat{\vec{y}} = f(\vec{w} \otimes \vec{x})$ **and its parameters** $\vec{w}$**, find** $\hat{\Gamma}$ **such that** $Err(\hat{\vec{y}}, \vec{y})$ **and** $Cmpx(\hat{\Gamma})$ **are minimum.**

$Cmpx(\hat{\Gamma})$ is some complexity metric operating on the dataset required to train the machine; be it simply the number of records/instances, entropy of the sample, or another user defined metric for complexity. A classic regularization term in machine learning reduces the complexity of the model, whereas $Cmpx(\hat{S})$ reduces the complexity of the sample.

The one-off nature of MT/EMT as defined above can be decomposed into a series of chunks that forms a timeseries of incremental MT/EMT (IMT/IEMT) to teach the machine incrementally.

*Definition 4 (***IEMT Problem***):* Given a model $\hat{\vec{y}} = f(\vec{w} \otimes \vec{x})$ and its parameters $\vec{w}$, find $\hat{\Gamma}_t$ such that $\lim_{\int_t^\infty \hat{\Gamma}_t}$, $Err(\hat{\vec{y}}, \vec{y})$ and $Cmpx(\hat{\Gamma})$ are minimum.

The $\int_t^\infty \hat{\Gamma}_t \rightarrow \hat{\Gamma}$ term is stating that as we accumulate the chunks through integration or summation in discrete cases, in the limit, the cumulative chunks should approach the overall sample $\hat{\Gamma}$ that is needed to approximate Γ.

Our definition of a curriculum above differs from Bengio et.al. [27]. In their case, a curriculum is seen as a weighting function on the sample with a higher weight initially given to a simpler problem to learn. Weights then shift to the more difficult problems in a similar manner to the classic Boosting algorithm [28], [29], where misclassified instances in one iteration receive higher weights in the following iteration, and an ensemble of learners get formed. Our methods, however, will generate modular structures that form a single network, while each module has a 'meaning' within an ontological representation.

Curriculum design assumes that Γ is known, while IEMT attempts to find ordered sets that in the limit will approximate $\hat{\Gamma}$. Formally, curriculum design partitions Γ such that $\bigcup_o \Gamma_o = \Gamma$. In IEMT, each $\hat{\Gamma}_t$ does not need to be a subset of Γ. These are only guiding lessons to lead the learner towards $\hat{\Gamma}$.

For a linear learner, each $\hat{\Gamma}_t$ results in $\vec{w}_t$. The objective of IEMT is to ensure that as $\int_t^\infty \hat{\Gamma}_t \rightarrow \hat{\Gamma}$, thus, from Definition 3, we draw $\lim_{t\rightarrow\infty} (\vec{w}_t) \rightarrow \vec{w}$. Each $\vec{w}_t$ defines a vector of angles, $\vec{\theta}_t$, with a dimension of $n - 1$.

So far, the only condition we have on $\hat{\Gamma}_t$ is that it is as short as possible. A second critical condition we need to establish is that $\hat{\Gamma}_t$ should be interpretable; that is, a meaning is associated with each chunk, and that the sequence of $\hat{\Gamma}_t$ over time

is, therefore, explainable by understanding the relationships among the meaning of the individual chunks and their order. This condition is what we call semantic ordering; that is, by sequencing chunks to the machine, with each chunk has a meaning that contributes to the overall logic of the learning process, we will have a lesson plan with identifiable rationale.

To achieve semantic ordering, we will need to design an ontology, Ont for the lessons. Each lesson needs to deliver a concept in that ontology, by generating the required sample of data to learn the concept. As the machine learner learns these lessons and compresses them using its internal representation, the process converges to a machine learner that has learnt the overall task, an explanation of how the learner has learnt the task, and a task-centred interpretation for each module. The above could be captured mathematically as shown in Equations 1- 2, where the $Linked$ operator synthesises the semantic order such that the synthesis covers the overall ontology.

$$Semantic(\hat{\Gamma}_t) \in Ont \tag{1}$$

$$Linked_t(Semantic(\hat{\Gamma}_t)) = Ont \tag{2}$$

IV. CASE STUDY

In this section, we present a synthetic case study using an artificial neural network to demonstrate scaffolding and the associated semantic representation. The 2D problem represents features calculated on different agents swarming in a space. A swarm is a group of agents that self-synchronise actions to achieve an effect or outcome. A slow swarm means that members of the swarm moves slowly, resulting on a longer time to reach the target; allowing them better opportunities to exchange information and learn from the environment. A fast swarm needs to get to its target fast, with lesser time to learn or exchange information.

Each cluster is represented by the mode of the time between successive visits of landmarks and the mode of the distance between the landmarks. The simple ontology is represented in Figure 3, whereby the primary concept of swarming is associated with two concepts. One is that a swarm has speed and the other is that a swarm occurs spatially over distance.
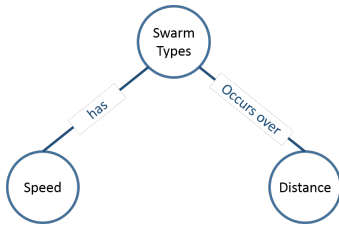

Fig. 3. Simple Swarm Attack Ontology.

In this dataset, the ground truth is represented with two inequalities given in Equation 3- 4. The rationale underpinning them is to define different forms of swarming. Equation 3 covers the space of slow swarming over medium ranges. Equation 4 covers the space of fast swarming over longer
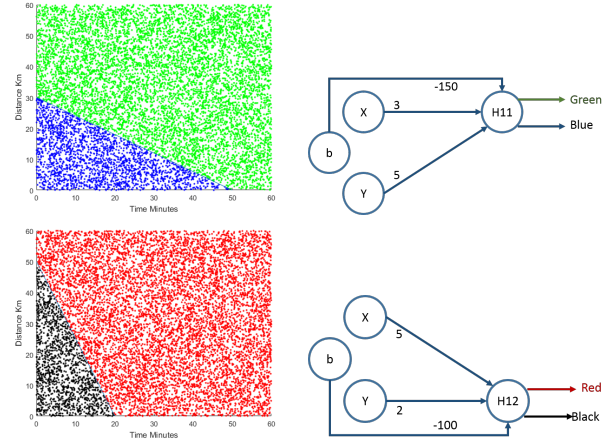

Fig. 4. Data Chunks.

distances. The intersection of the hyperplanes corresponding to these two inequalities create four polyhedra; three of which are polytopes.

$$3 \times Time + 5 \times Distance \geq 150 \tag{3}$$

$$5 \times Time + 2 \times Distance \geq 100 \tag{4}$$

If we consider this as the curriculum we wish to deliver to a machine learner, we could use the ontology to guide the sequencing of the datasets to the machine learner. Figure 4 presents the samples used to generate the two lessons and the resultant networks. The first two samples were generated to learn the lessons associated with Equations 3 and 4, respectively; signifying slow swarming over medium ranges, and fast swarming over longer distances, respectively.

As each concept is learnt by one network, these networks get frozen for the third lesson that synthesizes these concepts to learn the different types of swarming as shown in Figure 5.

V. DISCUSSION AND CONCLUSION

Machine teaching is an emerging field with a focus on the experience that gets presented to the machine, or what to learn. It complements machine learning, which focuses on how to learn. The literature is in its infancy, focusing on the data side with less or no emphasis being placed on the pedagogical framework for this educational process. In this paper, we brought concepts from education to introduce machine education, presenting to researchers and scientists the ingredients to educate machines in a similar manner to human education. System dynamics is proposed as an appropriate modelling framework to transform user intent and tasks into the skills expected from the machine and the learning goals. We then presented the concept of scaffolding, which, when coupled with ontology, offers a structured methodology to semantically decompose the problem into data chunks. Together with the system dynamic model, the approach allows the design of a transparent machine learners and offers a methodology that could systematically, experimentally or theoretically verify the
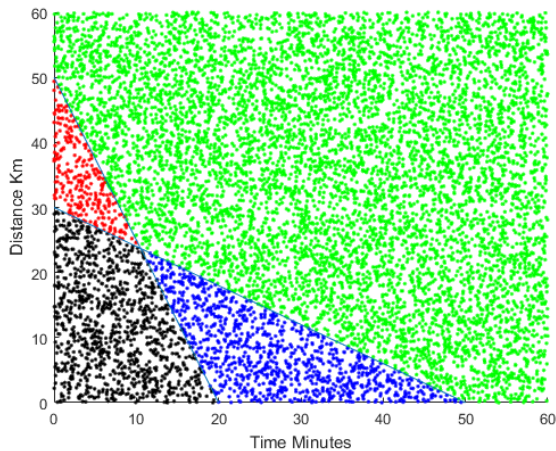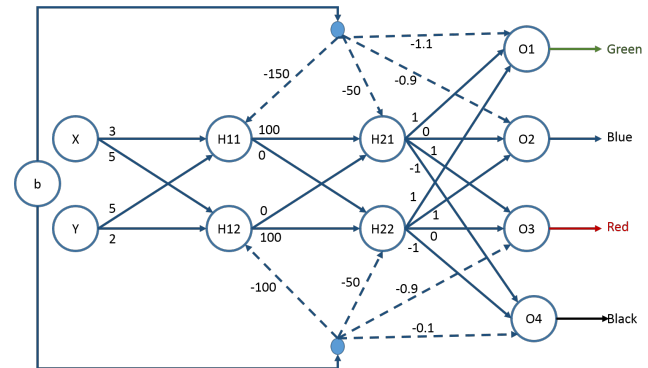
Fig. 5. Data Chunks.

learning experience of the machine to ensure its trustworthiness.

REFERENCES

[1] S. Mei and X. Zhu, "Using machine teaching to identify optimal training-set attacks on machine learners." in *AAAI*, 2015, pp. 2871–2877.

[2] B. Peng, J. MacGlashan, R. Loftin, M. L. Littman, D. L. Roberts, and M. E. Taylor, "Curriculum design for machine learners in sequential decision tasks," *IEEE Transactions on Emerging Topics in Computational Intelligence*, vol. 2, no. 4, pp. 268–277, 2018.

[3] G. Leu, E. Lakshika, J. Tang, K. Merrick, and M. Barlow, "Machine education - the way forward for achieving trust-enabled machine agents," in *NIPS'17 Workshop: Teaching Machines, Robots, and Humans*, 2017.

[4] X. Zhu, A. Singla, S. Zilles, and A. N. Rafferty, "An overview of machine teaching," *arXiv preprint arXiv:1801.05927*, 2018.

[5] N. R. Clayton and H. Abbass, "Machine teaching in hierarchical genetic reinforcement learning: Curriculum design of reward functions for swarm shepherding," in *IEEE Congress on Evolutionary Computation*. Wellington, New Zealand: IEEE, 2019.

[6] A. Gee and H. Abbass, "Transparent machine education of neural networks for swarm shepherding using curriculum design," in *International Joint Conference on Neural Networks*. Budapest, Hungary: IEEE, 2019.

[7] J. L. Elman, "Learning and development in neural networks: the importance of starting small," *Cognition Voume 48*, 1993.

[8] W. Dick, L. Carey, and J. Carey, *The Systematic Design of Instruction*. Merrill/Pearson, 2009.

[9] S. Dijkstra, F. Schott, N. Seel, R. D. Tennyson, and N. M. Seel, *Instructional Design: International Perspectives: Theory, Research, and ModelsVolume I*. Routledge, 2012.

[10] J. W. George, "Classical curriculum design," *A and H*, 2009.

[11] R. Jordan, S. Gray, M. Zellner, P. D. Glynn, A. Voinov, B. Hedelin, E. J. Sterling, K. Leong, L. S. Olabisi, K. Hubacek *et al.*, "Twelve questions for the participatory modeling community," *Earth's Future*, vol. 6, no. 8, pp. 1046–1057, 2018.

[12] A. Voinov, K. Jenni, S. Gray, N. Kolagani, P. D. Glynn, P. Bommel, C. Prell, M. Zellner, M. Paolisso, R. Jordan *et al.*, "Tools and methods in participatory modeling: Selecting the right tool for the job," *Environmental Modelling & Software*, vol. 109, pp. 232–255, 2018.

[13] S. Gray, A. Voinov, M. Paolisso, R. Jordan, T. BenDor, P. Bommel, P. Glynn, B. Hedelin, K. Hubacek, J. Introne *et al.*, "Purpose, processes, partnerships, and products: four Ps to advance participatory socio-environmental modeling," *Ecological Applications*, vol. 28, no. 1, pp. 46–61, 2018.

[14] S. M. Falconi and R. N. Palmer, "An interdisciplinary framework for participatory modeling design and evaluation – what makes models effective participatory decision tools?" *Water Resources Research*, vol. 53, no. 2, pp. 1625–1645, 2017.

[15] B. L. Barnhart, H. E. Golden, J. R. Kasprzyk, J. J. Pauer, C. E. Jones, K. A. Sawicz, N. Hoghooghi, M. Simon, R. B. McKane, P. M. Mayer *et al.*, "Embedding co-production and addressing uncertainty in watershed modeling decision-support tools: Successes and challenges," *Environmental Modelling & Software*, 2018.

[16] M. Duespohl and P. Döll, "Causal networks and scenarios: participatory strategy development for promoting renewable electricity generation," *Journal of Cleaner Production*, vol. 121, pp. 218–230, 2016.

[17] S. Elsawah, J. H. Guillaume, T. Filatova, J. Rook, and A. J. Jakeman, "A methodology for eliciting, representing, and analysing stakeholder knowledge for decision making on complex socio-ecological systems: From cognitive maps to agent-based models," *Journal of environmental management*, vol. 151, pp. 500–516, 2015.

[18] S. Elsawah, S. A. Pierce, S. H. Hamilton, H. Van Delden, D. Haase, A. Elmahdi, and A. J. Jakeman, "An overview of the system dynamics process for integrated modelling of socio-ecological systems: Lessons on good modelling practice from five case studies," *Environmental Modelling & Software*, vol. 93, pp. 127–145, 2017.

[19] A. Voinov and F. Bousquet, "Modelling with stakeholders," *Environmental Modelling & Software*, vol. 25, no. 11, pp. 1268–1281, 2010.

[20] S. Elsawah, A. McLucas, and J. Mazanov, "An empirical investigation into the learning effects of management flight simulators: A mental models approach," *European Journal of Operational Research*, vol. 259, no. 1, pp. 262–272, 2017.

[21] R. J. Scott, R. Y. Cavana, and D. Cameron, "Recent evidence on the effectiveness of group model building," *European Journal of Operational Research*, vol. 249, no. 3, pp. 908–918, 2016.

[22] J. D. Sterman, "Learning in and about complex systems," *System Dynamics Review*, vol. 10, no. 2-3, pp. 291–330, 1994.

[23] ——, *Business dynamics: systems thinking and modeling for a complex world*, 2000, no. HD30. 2 S7835 2000.

[24] J. Hammond, P. Gibbons *et al.*, "Putting scaffolding to work: The contribution of scaffolding in articulating esl education," *Prospect*, vol. 20, no. 1, pp. 6–30, 2005.

[25] R. A. Fisher, "On the mathematical foundations of theoretical statistics," *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, vol. 222, pp. 309–368, 1922.

[26] A. Barron, J. Rissanen, and B. Yu, "The minimum description length principle in coding and modeling," *IEEE Transactions on Information Theory*, vol. 44, no. 6, pp. 2743–2760, 1998.

[27] Y. Bengio, "Learning deep architectures for AI," *Foundations and trends® in Machine Learning*, vol. 2, no. 1, pp. 1–127, 2009.

[28] R. E. Schapire, "The strength of weak learnability," *Machine learning*, vol. 5, no. 2, pp. 197–227, 1990.

[29] Y. Freund and R. E. Schapire, "A desicion-theoretic generalization of on-line learning and an application to boosting," in *European conference on computational learning theory*. Springer, 1995, pp. 23–37.