BeliefNet: A neurosymbolic model to enhance context based traversability predictions for autonomous agents in complex environments

Tom Scott, Argyrios Zolotas and Yang Xing
Centre for Autonomous and Cyber-Physical Systems, Cranfield University, UK

Abstract. Knowing how to traverse complex unstructured environments is a difficult and multivariate challenge, but one which humans can achieve through logic, reasoning and experience, yet some of the most beneficial use-cases for autonomous systems, require them to operate effectively in complex environments without regular or significant human intervention. Furthermore, for machines to support humans in some of the more critical use-cases trust in decision making will be crucial, ensuring operators have confidence to deploy the capabilities. Inspite of its importance, enabling autonomous agents to navigate effectively and reliably in complex terrain is a difficult and unsolved challenge. Advances in neurosymbolic AI present an opportunity to significantly enhance performance in complex, explainable and uncertain decision making, such as autonomous traversability analysis, by drawing together symbolic reasoning with the learning capability of neural networks. The challenge of complex environments is complicated by its non-deterministic nature, terrain will adapt and change through domains and its properties can adapt rapidly based on external factors like weather, or objects that are in proximity, what is true for one location on one day, will not persist. This paper presents a new neurosymbolic model structure that was designed specifically for this task. It uses experience to build a world model, similar to that of a neural network, but with some key delineating features, such as, full explainability, through life adaption or evolution and zero-shot capability, enabling it to perform as both a reasoning engine and a memory representation for an autonomous system. This provides the reasoning backbone for an autonomous agent to determine the level of risk each object presents based on its context and therefore determine the best possible route.

Keywords: Neurosymbolic AI, Machine learning, Knowledge based learning, Autonomous systems, Complex environments

1. Introduction

Autonomous systems present opportunity to transform how humans complete some of the most dangerous, unpleasant or persistent tasks. This is especially relevant to domains such as Defence or Search and Rescue, where autonomous systems can remove humans from hazardous situations. These use-cases present some of the greatest beneficiaries of autonomous systems, but have some of the most demanding requirements, most notably the ability to operate reliably in very complex terrain and dynamic domains, whilst maintaining a high degree of trust by their operators. Robustly operating in complex environments requires platforms to operate in both unstructured and uncertain terrain, where clear transition points between features may not exist, with high variation in slope, roughness and unpredictable terrain features like holes or depressions [51], [54], [55]. Furthermore the characteristics of an object cannot be determined effectively without understanding the context in which it is found. To navigate, it requires inductive and deductive reasoning, an understanding of the environmental conditions, probabilistic judgment, domain adaption and the ability to handle uncertainty. Whilst humans can thrive in these situations, when considering autonomous agents neither a symbolic or neural approach replicates them all sufficiently. Neural approaches generally fail to reason effectively and suffer from a lack of explainability, but can be adaptive to out of distribution data, whilst symbolic approaches can reason but require a significant upfront knowledge base and cannot effectively generalise. Fundamentally performing these activities within an autonomous platform is not a simple extrapolation of either approach. It is for this reason that this remains an outstanding challenge in the field of autonomous systems [15], to the degree that in 2021, it was released as a key area of research by the US research agency, DARPA within their RACER programme [66]. Despite its inherent challenge, it is a key determinant of the application of autonomous systems to the use-cases within the Defence domain, with the UK establishing two organisations dedicated to furthering its capability [58] and a long term strategy to deliver against the use-cases.

Humans are able to operate dynamically and effectively in complex terrain, making risk based decisions and generalising from their experiences, enabling them to predict how previously unseen terrain may react. Yann LeCun recently outlined that to enabling machine autonomy we may need to enable systems to replicate humans and, most notably, our ability to generate a 'world model' [31]. LeCun takes the perspective of considering how the human brain performs, rather than how it is structured, this is not revolutionary, and was in fact was a founding component of cybernetics [45], but has been somewhat overshadowed by the development of models such as the transformer architecture [18]. Creating a more 'human-like' thought process for autonomous systems has a number of benefits, outside any direct performance gains, such as transparency, explainability and trust. Tools such as reasoning and logic are easier to interpret than the complex mathematical functions of deep learning models [61], which in themselves are crucial for adoption into high risk and highly regulates use-cases. By definition, a reasoned or logical argument must transcend the originator, it must be communicable and be able to be understood externally, even if not always agreed with [37], therefore, an operator could interpret the reasoning behind a decisions, even if the decision was wrong, understanding why it was wrong, makes it a deterministic action and vastly increases trust [22]. This paper presents a model architecture that looks to establish a world model, and enable transparent and reasonable decison making in complex environments.

The challenge of navigating in complex terrain is not a single problem to solve and can be considered across a number of areas, such as perception, localisation, cognition and motion control [43]. Whilst each has their challenges, this paper focuses on cognition, and specifically on how to enable an agent to determine the traversability of an object by learning how it adapts to the context and environment in which it finds itself. This concept paper builds upon the world model concept, using neurosymbolic AI [26] and drawing upon more traditional fields such as epistemology, to develop a human-like approach to solving the autonomy in complex environments challenge. Through inductive learning, deductive generalisation and the agents application of adaptive beliefs, it enables the platform to build a world model enabling it to deduce and predict the traversability of a given object, based on its situational context and evolve its beliefs for new domains.

Mileo et al define one of the key goals of neurosymbolic AI in explainability is to achieve 'neural cognitive mapping' deriving high level concepts, relations and reasoning from low level data [39]. This paper contributes towards this objective through the presentation of a new model structure, BeliefNet and assessing its performance in the autonomous systems domain. BeliefNet uses a symbolically built neural architecture to form experience based beliefs, representing causal relationships between a target object, its context and traversability risk. The model can evolve continually with an agent's experience and requires small amounts of data to learn, enabling the agent to adapt quickly to new environments, whilst its symbolic information ensures the evolution remains deterministic. The model seeks to extrapolate causal relationships, enabling it to generalise effectively to different domains. It also enables a belief based inference, in which predictions are deterministic and explainable, enabling operator trust.

The contributions of this paper are as follows:

– The proposal of a new Neuro[Symbolic] model structure, the BeliefNet.
– A demonstration of the BeliefNet performance on an adapted version of the Yamaha CMU [62] dataset, in increasing agent cognition performance in complex environments.
– Comparison of context based terrain traversability prediction and object based prediction.
– A high level traversability taxonomy for ground platforms based on risk and speed.

1.1. Defining the challenge

Before further outlining the approach, it is important to first clarify the specific challenge and the constraints placed on an agent within the given situation. To bound the use-case we will use an example from a previously trialed experiment within the Defence domain. This saw autonomous vehicles being used to deliver supplies to forward deployed units, it involved traversal of challenging off-route and dynamic terrain [44]. This use case presents a number of constraints that we must consider. The terrain is extreme off route, it is likely to include complex objects such as tall grass, and present no simple, viable option, this is likely to inhibit the use of tools like Lidar [63]. The task is of reasonable length with the agent expected to operate between 5km and 10km without human intervention, the agent is likely to be required to make uncertain decisions. The agent is likely to have not previously seen the terrain before and data sample size that reflect the exact environmental conditions are likely to be small, inhibiting conventional machine learning. Finally the task requires high levels of user trust. The operator will need to be confident in the agent's ability to achieve the desired objective with high confidence, before being willing to initiate the task.

This poses a number of key questions, which are represent the context for BeliefNet:

1. How can the agent build a world models, when access to large quantities of representative data may not be available?
2. How can the agent take what it knows and adapt it to a new unseen environment?
3. How can the user better trust that the agent will complete its task?
4. How can the agent use vision based systems to accurately assess traversability?

2. Existing work

When considering the existing work in this field, we must consider two areas as this paper proposes a solution to the use-case, and what is believed to be a new neurosymbolic architecture. Firstly the existing approaches to enabling autonomous cognition in complex environments. Secondly, the field of neurosymbolic AI more broadly.

2.1. Approaches to agent cognition

The field of traversability has had a great deal of focus in recent years, leading to three primary approaches emerging to conduct traversability assessment, terrain classification, terrain mapping and end-to-end [5].

Lidar analysis has been used extensively in traversability assessment approaches, both in direct obstacle avoidance [57][30], or in more complex feature segmentation [20] [2] [64]. Whilst delivering promising results such approaches are spatial by nature, potentially over simplifying the traversability calculus by ignoring the environmental and situational semantics. Furthermore the active nature of Lidar presents challenges in use-cases where the light emissions have a negative secondary effect.

Terrain classification presents a method of making use of semantics, due to significant advances in computer vision, with the introduction of models such as YOLO [48] and approaches like vision transformers [14] and panoptic segmentation [65], this has become increasingly, enabling real-time inference on edge based devices. The use of computer vision enables terrains to be segmented into constituent objects, which can then each be evaluated based on their characteristics. As the terrain classification of complex environments is non-trivial, due to the discontinuous nature of objects, feature overlap and environmental conditions [34], this continues to be an area of active research [11], [16], [21], [10], [59] [2], [55], [17]. Though there has been significant research into the visual modality, vision and Lidar have been combined to enable semantics to be integrated with the spatial data of Lidar [36]. Each of these has taken different approaches in an attempt to capture the nuance and complexity of the off-route environment. There are two components to the challenge, first detecting and separating a specific object, then assessing the traversability of the object. Though some, like [2] have integrated both components, most parts of the research primarily focus on accurately determining the object, not assessing the traversability. One challenge in this approach is that it can neglect the need to consider the environment and context for a specific object, there will be some objects which will directly impact the traversability of others. Without this context it can be challenging to make an accurate reasoned determination.

End-to-end deep learning approaches have had success in classifying the traversability of an image [60], [32]. Self-supervised approaches have been used in this area, in which a platform trains a model based on its own experiences to predict the traversability of the terrain [50] [52], to negate some of the challenges around labelled data availability and data variance. Such approaches however can be limited in generalisation performance and crucially has limited explainability due to the conventional neural architecture of end-to-end deep learning approaches. Furthermore, using a single model to assess traversability could fall into the trap of ignoring a platforms performance characteristics, which are likely to be very different for separate types, models or even age of a platform.

In a more conceptual area of the field, human cognition has been an inspiration for machines, from Walter's Tortoises developed in the early era of cybernetics [45]. The Cognitive Patterns Knowledge Generation system, which used pattern generation, extracted from long term memory and integrated into a perception module as a method of enabling agent reasoning to generalised cases [56]. More recently LeCun proposed his methodology for enabling reasoning using the concept of a world model which more accurately represented human thought, than conventional deep neural networks [31]. These all focus on attempting to solve a key problem, enabling an autonomous agent to reason using their experiences in unseen, complex or unpredictable environments.

2.2. neurosymbolic AI

A promising area of research for integrating reasoning within the field of AI, is neurosymbolic AI, which suggests to provide advances in machine decision making, representation, explainability and reasoning [7]. This area of research seeks to identify ways to integrate the benefits of symbolic reasoning the generalisation and learning power of sub-symbolic, connectionist or neural network approaches [26], building upon the system 1/system 2 approach defined by Kahneman [25]. The field is still growing and there remains alot of diversity in approaches, but all have in common the structure of perception, integrated with existing knowledge, focusing on increased reasoning and explainability [53], and are particularly applicable to use cases with high levels of human-machine interaction [4]. [13] outlines six architectural approaches to neurosymbolic AI, building on those initially defined by Kautz. Two of particular relevance to this paper are defined as NEURO;SYMBOLIC and NEURO[SYMBOLIC] [28].

NEURO;SYMBOLIC reflects a system type where a symbolic and neural system work in concert with each other, communicating and passing information between them, to achieve a common objective [26]. Examples of this are knowledge graph integration with neural networks [13] which enable a neural network to query from, input to and be validated by symbolic knowledge graphs. NSNnet, which passes between neural and symbolic modules in an aim to solve hand written sodoku challenges, presents a unique perspective that maps both input and output to a non-symbolic output, with a central symbolic reasoning engine [1]. Both these examples are dependent on a core level of symbolic reasoning. The Neuro Symbolic Concept Learner (NCSL) designed to unify text and visual concepts through learning image and question-answer pairs [35]. This model presents an interesting advance as it enables symbolic concepts to be learnt, without implicit knowledge being defined upfront.

In contrast the NEURO[SYMBOLIC] system is one in which a neural network learns to reason about relationships between neural entities [26] [29], in effect forming a neural network of symbolic entities. This is perhaps the most complex and least mature of the areas of research within the field. Logic Tensor Networks (LTN) and Logic Neural Networks (LNN), which form networks from symbolic relationships and enable weighted training of the relationship using back propagation based on a set of first order logic statements [3] [49]. The pLogicNet model mostly precedes the core definitions of the neurosymbolic AI represents a method similar to the Logic Tensor Networks based on the application of Markov Logic Networks [47]. The LTN and pLogicNet are designed to improve, validate or deconflict a set of apriori logical statements. The challenge with these approaches when applied to an agent based approach is that it requires the upfront knowledge which may not be practical to achieve. Models such as the Neuro Symbolic Reinforcement Learner, INSIGHT, by Luo et al uses a neural network to learn symbolic policies supporting the agent in its decision making enabling reasoning to be learnt from the environment [33].

The current state of neurosymbolic AI presents significant advances in both reasoning and explainability, the NEURO[SYMBOLIC] concept of an single network encapsulating symbolic reasoning, presents an opportunity to be a world model for an agent. Requirements for upfront knowledge however limit the ability for the model to learn and adapt through-life, limiting the ability for the model to adapt. BeliefNet has taken the concept of a symbolic network trained using sub-symbolic approaches, but in a manner that reflects the domain learning capabilities of models such as INSIGHT or the NSCL to enable adaptive reasoning for an agent operating in a multitude of complex terrains.

3. Context based classification

To understand the benefit of the BeliefNet model approach, it is first important to understand the fundamental strategy which underpins the traversability prediction, context based classification. We can consider there to be two types of object, ones with an immutable traversability value, such as a tree or a fence, the value remains constant regardless of the scenario, then objects with mutable values, in which the traversability value is dictated by more than just the object class alone. The premise is that for a given subject mutable object, for which we are seeking to make a prediction, the traversability risk value is a function of a number of external context drivers, such as the presence of surrounding objects, the weather (current and historical) and the domain in which it is in. When combined these factors can provide a more nuanced and accurate prediction of traversability risk. The images at 1 show an example of two separate examples of the same trail object, which will both react differently based upon the surrounding context. Determining how a given context will impact an object is non-deterministic and highly dependent on the base object, some such as complex vegetation or soil can have huge variance, where as objects such as low grass or hard trail can have a far smaller variance. The number of combinations, number of context objects and the naunced relationships that exist between them, make understanding the impact of context a complex reasoning task, least of all because in complex environments, there is not a finite set of objects that may be part of the context. It is this specific type of complex reasoning in which the BeliefNet architecture has been designed to perform.

4. Model approach

BeliefNet at a high level abstraction appears to reflect the overall structure of a neural-network, and it has been actively designed with this intent. It does however have a number of key functional divergences towards an integrated neurosymbolic structure, it is these separations which have enabled it to adapt to the complex task of risk with very small amounts of data when compared with a conventional neural network, whilst retaining absolute explainability in the models deduction. The core premise of the model has its founding in epistemology, most notably the concept of beliefs and knowledge. Beliefs are something which the system expects to be true [42], with varying degrees of confidence based on its own experiences, and knowledge being something that extends a belief to be justified and true [68]. As humans our beliefs rarely remain static, continually updating and adapting to our experiences and our current domain, this opposes a traditional neural-network approach. When we face something unknown, we find the set of closest beliefs, use them to make a prediction, then create a new belief which captures the separation between the prediction and truth, often captured within the concept of Predictive Coding [40]. It is this function that the BeliefNet model looks to model, whereas a conventional neural-network is expected to remain constant once trained, BeliefNet is designed in a manner that makes domain adaption and updating a core function of the model.

The BeliefNet model is designed to operate post perception, so can be agnostic to the object classification model, or even the modality, it only requires a common ontology with the classification model, and takes a set of symbolic predicates as an input. Its objective is to generate an optimal set of beliefs, inferred from the training data, which can be dynamically combined to form accurate predictions in achieving an overall goal, in this instance, the goal is predicting the correct traversability value. The model is logically build, meaning connections are only generated between objects that have been seen together, it is not fully connected at any stage. Combined with the symbolism retained within each node, it affords the ability to only activate small portions of the model using inference, which aids explainability, but also provides the model with its reasoning capability when facing unknown situations. The ability to handle this uncertainty is crucial due to the high likelihood of out of distribution data appearing, with

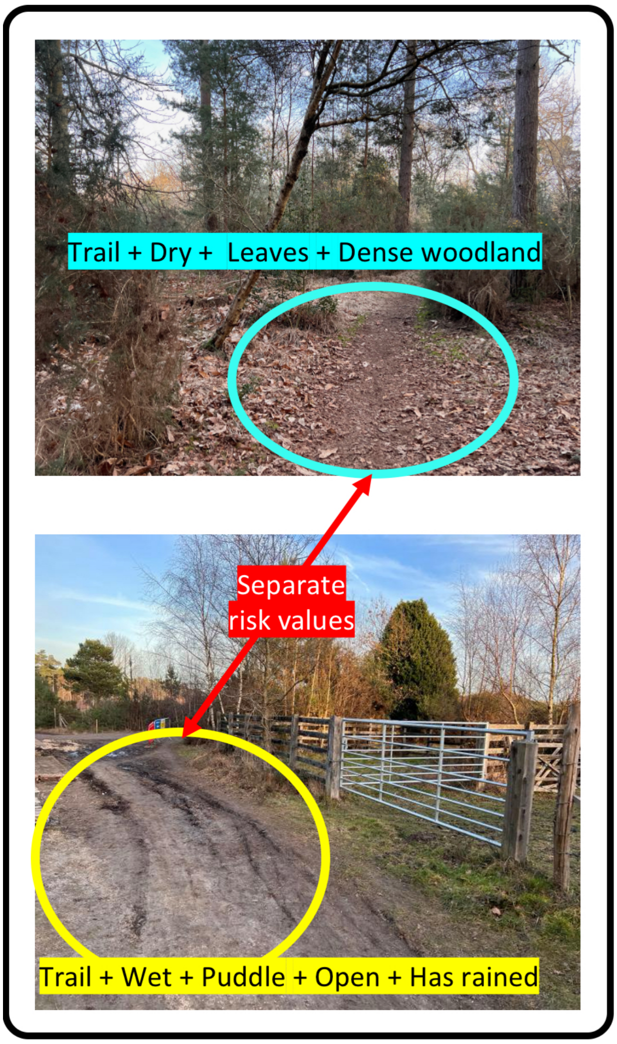

Fig. 1. An example of two instances of trail objects with their surrounding context, resulting in separate risk assessments

regularity during inference in complex environments. This approach acts as a zero-shot domain adaption model, without the need for the high data volumes conventionally required through existing zero-shot approaches.

The model incorporates the concept of belief and knowledge. Knowledge holds a higher weighting within the model and can be fixed or trained as part of the optimisation. Whilst beliefs hold randomly initialised weights and biases, which adapt with evidence and justification. As mentioned previously, the model uses a lazy-relationship model, in which node relationships are only formed at the point they have been identified, this includes with output nodes. As a result if only a single output has been observed from a given combination of predicates, there will only be a single output value returned, in effect representing learnt knowledge. At the point this diverges, and a second output node is generated it represents a belief with the weights to each output node representing the level of confidence in each output. For example, the system may have never have observed grass with snow and a puddle before, but it has observed grass with snow and grass with a puddle, reasoning over these beliefs can be used to estimate how grass is traversed. The model uses the belief and knowledge structure as a foundation for human machine teaming. Operators can increase or decrease belief confidences directly through manual training iterations. Secondly, operators can specify specific knowledge into the model directly.

5. Model structure

The model is formed of a number of components, some which are adaptions of existing deep-learning approaches and some which are specific to BeliefNet. At a high level the model should be considered as post-processing of a perception model, it initialises by taking the perception output and transforming this into a graph structure using a common ontology, known as an instance graph. The instance graph is generated in the complex environment use-case as the output of a semantic segmentation model, such as YOLO [24]. The predictions are further enhanced through a depth perception model [6], estimated 3D separation between objects, and augmented with environmental tags, representing the weather, light and domain. The instance graph is a dense symbolic representation of a given image. During training the instance graphs are converted to a series of context graphs, representing the surrounding objects, distances and environmental tags for a given object for which a prediction is made. Context graphs are passed to the building algorithm, which is a custom training method designed to extract causal relationships between objects, context and a traversability value. The output of the build is a partially connected model architecture, which is an inductive representation of the beliefs and knowledge extracted from the training data. This forms the basis of the neural network, each node holds an activation function (which can be varied across the model), bias parameter and each edge has a connection weight. This enables the second phase of training, which uses conventional back-propagation using the Adam algorithm [27] to adapt the weights to in a supervised manner. Once complete this forms the base BeliefNet, from this inferences can be made, and crucially, additional nodes can be formed and adapted as the agent learns of new information. It is the logical build process before the back-propagation which provides the reasoning capacity and explainability of the structure.

The model is fundamentally built on a graph structure, with nodes representing a logical predicate and relationships representing the weight parameter. The nodes are either an input, belief or output node, each is symbolically represented in the structure. Conversely to most networks, BeliefNet depends on referencing the symbolic name of each node, input nodes hold atomic predicates, such as an individual object or context variable. Belief nodes are the combination of their related predicates, in effect an AND node, which are again referenced as a symbolic statement. Output nodes represent a given output value for a given prediction object. A single model may have multiple output layers representing separate objects, enabling knowledge transfer and generalisation. It is the symbolic structure of the nodes, combined with the logical build that enables the model to retain the lineage and attribution of outputs, affording it a high degree of explainability. The architecture at 2 shows how these components fit together within the model.

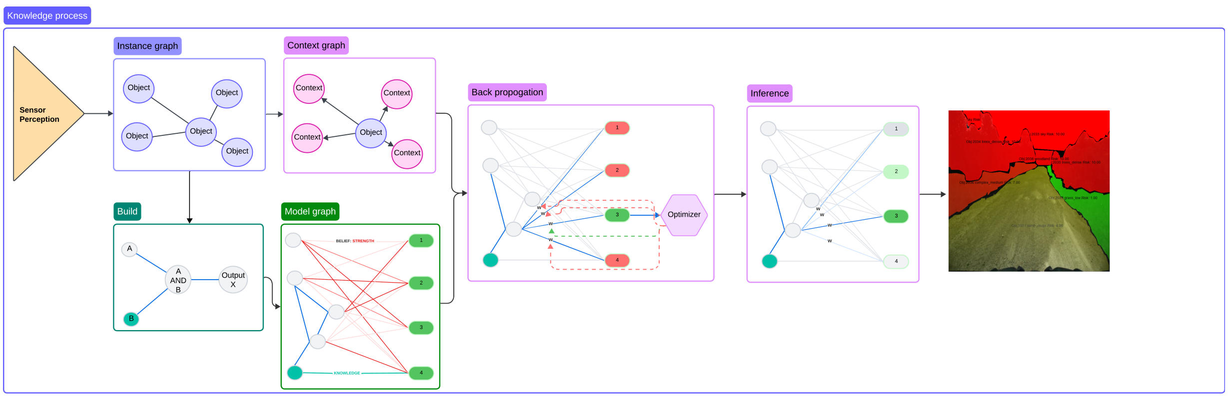

Fig. 2. Model Architecture, the high level architecture of the model is based on the structure of a neural network, but with adaptions to enable the symbolism to be retained throughout training and inference.

5.1. Key components

5.1.1. Belief goal

The model is designed to define a set of beliefs and accompanying weights and biases to maximise prediction to achieve a specific learning goal. The goal is fundamental to the models ability to learn and defining the goal will have profound impacts on the models ability to learn. In the case of this use case the goal is the prediction of a traversability index value for a given object and context. A traversability value could be infinitely complex and is very specific to an individual agents performance characteristics, as such we sought to define the goal to a level of abstraction where it could be generalised across platforms and interpreted for their needs. Furthermore, for the model to operate effectively, the goal values are required to be discrete not continuous. This led to the development of a traversability index. The traversability index categorises expected speed (relative to an agents default) and the level of caution the agent will require in their traversability. The traversability risk analysis framework proposed by [15], in which multiple metrics such as risk or collision, slippage and contact-loss are combined into a single measure of risk, as the basis for a unitary caution value. Whilst it can be common to see traversability risk as a regression problem [23], the platform considerations can be abstracted into discrete categories, creating additional fidelity and evolving the problem into one of classification. This enabled 11 distinct values to be defined, which were at a level of abstraction which meant relative traversability could be compared, whilst enabling an agent to generalise the values to their performance characteristics. These values are shown in the diagram at Figure 3. As the model represents a number of prediction objects, there is a specific output layer for each type, which can be symbolically referenced. This output structure provides benefits in generalisation, enabling outputs to be assessed if context has not been seen in training. This approach also sets the foundation for cross-task generalisation, in which separate layers can exist for multiple tasks. Currently it uses for object traversability risk layers, however this could be more granular, with layers for variables like speed, roughness and traction, each using the common model backbone.

5.1.2. Context graph

Within a given instance, there may be multiple objects about which a traversability assessment may want to be made. For each of these, a context graph (G) is generated, representing all the objects (V) with relationships (E) to the target. Captured as a sub-graph of the overall instance, it captures the target object(t), context object(c), relationship type (r) and the strength(s). For the traversability use case the relationship is the positional relationship of the two objects, and strength represents the 3-dimensional euclidean distance which is generated as post-processing from semantic segmentation. To ensure that this remains a sub-graph, a relationship threshold is set. The relationship threshold, and category ranges are parameters which can be tuned within the model.

$$G = (V, E) \tag{1}$$

Each edge $e \in E$ is defined as:

$$e = (t, c, r, s) \tag{2}$$

5.1.3. The model

The model M can be represented as a combination of nodes N and edges E, in line with any conventional graph. However, there are a number of node types within the model, each characterising different behaviours, input node I, belief node B and output node Φ such that $N = I \cup B \cup \Phi$. The edges are directional relationships between two nodes, a predicate and a logical relationship, $n \in N$. Each edge hosts a trainable parameter, representing the weight, each node represents a trainable bias parameter. The model input layer represents all possible atomic predicates, output layers are three-dimensional, with separate output layer existing for each prediction object or prediction task, in between exists the beliefs.

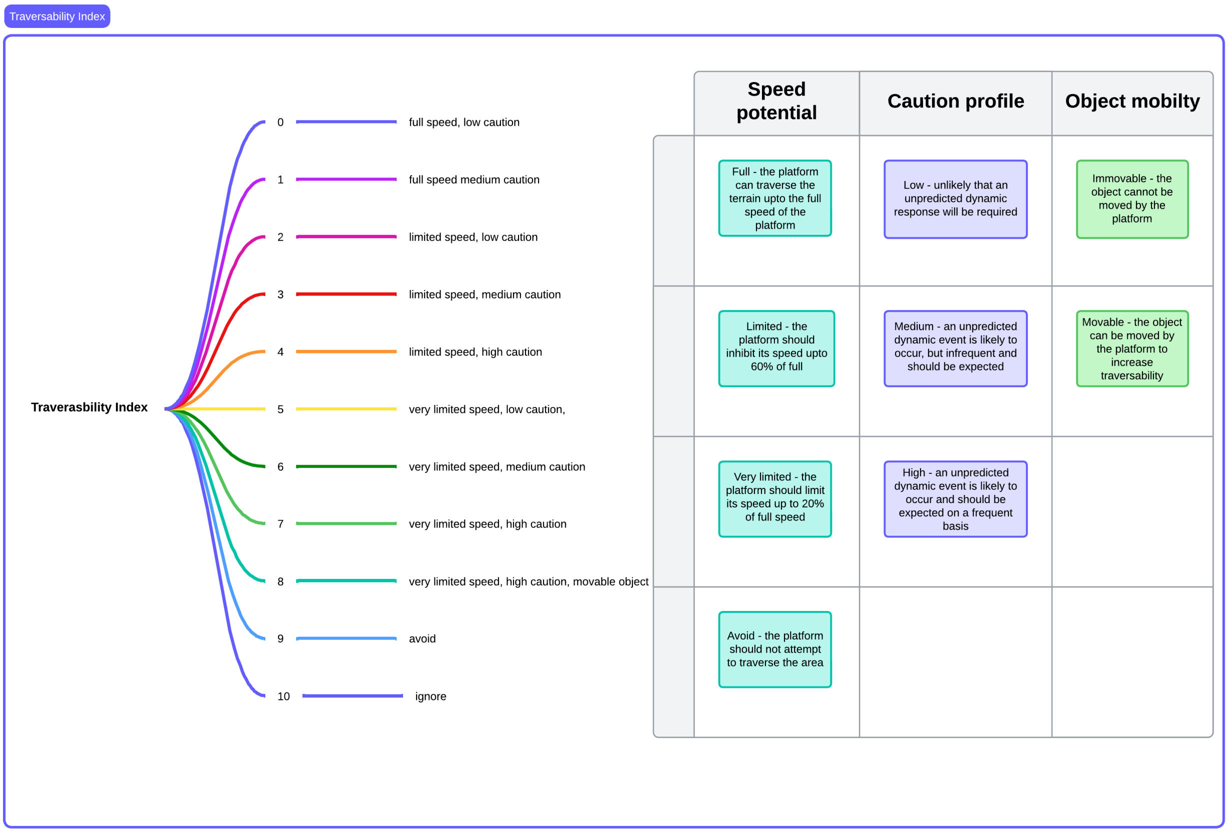

Fig. 3. Traversability Index, there are 11 discrete traversability components, which increase in complexity as defined by the variables in the right hand table. These are categories that dictate the relative speed, level of caution the platform requires and the mobility of an object. They are assessed based on the individual perception of a single platform, therefore these cam be considered relative to the performance characteristics of an specific platform.

5.1.4. A Belief node

A belief node provides the foundational component of the BeliefNet model, loosely equivalent to a single neuron within a neural network.

$$
\begin{aligned}
& b \in B \\
& C = \{c_1, c_2, \ldots, c_n\} \\
& \phi_t = \{\phi_1, \phi_2, \ldots, \phi_n\} \text{ where } \phi_t \in \Phi
\end{aligned}
\tag{3}
$$

Where (B) is the list of all beliefs within the network, (C) is a subset of context objects that appear simultaneously, this can be only one and (ϕ_t) is the set of possible traversability outcomes for the target (t).

As with a conventional neuron, each $c \in C$ has both an input value, $i_c \in I$ and a weight $w_c \in W$, there is a bias term b and an activation function to account for non-linearity, *act*. Meaning the output value of a belief node is [46]:

$$
o_b = act((\sum I * W) + b) \tag{4}
$$

The belief node can hold multiple output relationships, during the build, it is initially formed with a single relationship, containing a randomly initialised weight, $w(b_n, \phi_{tn})$. As the build continues additional weights are generated as required to additional components of ϕ_t or to other belief nodes, $w(b_n, b_m)$.

Belief Nodes form the basis of the symbolic nature of the model, the foundation of this is the node naming structure, which is based on the predicates. When combined into belief nodes, predicates and their logical relationship are retained within the name of the node. This enables individual nodes to be referenced directly and for the contributing predicates to be directly identified.

$$\begin{aligned}
&\text{Let } p_1 = grass_low, \quad p_2 = hardcore_smooth, \quad \theta = newbeliefnode \\
&\text{If } \langle p_1, e, \theta \rangle \text{ created and } \langle p_2, e, \theta \rangle \text{ created,} \\
&\text{Then } \theta = p_1 \wedge p_2
\end{aligned} \tag{5}$$

Output nodes are similarly a subclass of belief nodes, which represent a specific output categorisation. Nodes are combined into layers, in which each node represents a traversability index value, and a the layer is indexed to the object being classified. This provides the model the ability to classify multiple different objects with the same model backbone. As they are a multi-class classification output, each output layer is combined with a Softmax function [9]. It is important to note the Softmax only applies to the specific prediction object output layer, not all outputs. Input nodes are a subclass of belief nodes which represent a single atomic predicate. During the build phase each $c \in C$ is represented by an input node. The input to the input node can be adapted, but performs best when represented by the relationship strength within the context graph, s, demonstrating the distance between objects. They often have separate activation functions than a general belief node within them due to the single input value of the node. Input nodes are characterised as:

$$o_i = act(I + b) \tag{6}$$

5.1.5. Model build process

The model is designed to be persistent and adaptive throughout the lifecycle of an autonomous agent, meaning it can be trained from no beliefs or use new instances, gained through experience, to update existing beliefs, both use the same build methodology. Conventionally neural networks have an initialised architecture, which remains constant throughout the life-cycle of the model, this is central to the matrix multiplication approach the model uses. However, this inhibits adaptability and the ability to integrate new knowledge when the model is not fully connected, integrate new predicates or determine the cause of a given output value. As a result BeliefNet integrates a build phase prior to model training, in which relationships between predicates, beliefs and outputs are dynamically formed, based upon presence in the training set. As it does not rely upon matrix multiplication to generate an output, once built, new relationships, new predicates and crucially new beliefs can be added dynamically, meaning it evolves with the agents understanding of the world.

The build process uses individual instance graphs, where each given instance is a component of the overall training set of instances $ins_m \in Ins_{\text{train}}$. Each instance is formed of a set of context objects (C^{ins_m}) and a set of target labels $\{\phi_{tn}^{(\text{ins}_m)}\}$, when the target label is a subject object and value, representing the traversability of the object being made a prediction about. The model first establishes that each $\phi_{tn}^{\text{ins}_m} \in \Phi$ and that all $c_x^{ins_m} \in I$, else new predicate nodes are created. It then seeks to identify an existing belief which matches the exact context where $b_y \equiv C^{\text{ins}_m}$, if found it confirms that $\phi_{tn}^{\text{ins}_m} \in W_{b_y}$, else it creates a new relationship $w_{b_y}, \phi_{tn}^{\text{ins}_m}$. If no direct match is found, the function searches for the existing beliefs which host partial matches such that $b_y \subset C^{\text{ins}_m}$, it then creates a new belief b_z formed of $w(b_y, b_z)$ and $w(C', b_z)$ where $C' = C^{\text{ins}_m} \setminus b_y, b_z$. If no partial matching beliefs are found, it creates the belief from the relevant input nodes directly. This process continues through each instance in Ins_{train}, or can be replicated when a new instance is identified. For each relationship, the parameters are randomly initialised to prevent biasing the model into a local minimum. This can be demonstrated in Algorithm 1:

Algorithm 1 Node and Relationship Update Procedure

Require: Features $F = \{grass\_low, puddle, tree\}$, Label $L = 1$

Let $N \leftarrow grass\_low \wedge puddle \wedge tree$
if Exists(N) **then**
    if Exists(Relation(N, e, L)) **then**
        Do nothing
    else
        CreateRelation(N, e, L)
    end if
else
    Let $P \leftarrow grass\_low \wedge puddle$
    if Exists(P) **then**
        Create($N = P \wedge tree$)
    else
        Create($P = grass\_low \wedge puddle$)
        Create($N = P \wedge tree$)
    end if
    CreateRelation(N, e, L)
end if

The model build can be augmented with a-priori knowledge during the build phase, where testimonial knowledge can be represented in effect in first order logic. Relationships between specific predicates can be unified as knowledge with a direct relationship to the output node. This alone would not be sufficient to capture knowledge, it is therefore knowledge nodes are initiated with high default parameter values for the weights and biases, often 1, this value has obvious impact on the model, so the value must be tested based on the domain. These parameters can be included or excluded from the optimiser, meaning they can be fixed or adapt with back-propagation. This represents the fact that knowledge could be permanently infallible, which is useful for human defined 'red-lines', or could be feasibly disproved by future evidence. Both are viable options within the model. This feature enables the model to draw on some of the benefits of tools like the LTN [3], which reasons over a corpus of provided knowledge, whilst enabling the system to add or adapt this knowledge based on induction. This is core to the domain adaption capabilities of the model. A knowledge node could be represented as follows:

Knowledge predicates are defined:

$$KP = \{c_1, c_2, \ldots, c_n\}$$

A relationship is formed between the knowledge predicates and the output node:

$$\langle KP, e_k, \phi_{t,n} \rangle \tag{7}$$

The relationship weight is dictated by the knowledge type:

$$w_{e_k} = \begin{cases} w_m, & \text{if knowledge is mutable, trainable} = \text{True} \\ w_n, & \text{if knowledge is immutable, trainable} = \text{False} \end{cases} \quad \text{where } w_m < w_n$$

5.1.6. Dynamic activation

The concept of relevant beliefs is also a separation from conventional ML, but one which has been seen in neurosymbolic AI through the freezing of specific input nodes and network dissection [38]. The input layer is considered to be all atomic beliefs (those of the lowest fidelity) from a given context graph, only the atomic beliefs represented in the graph are activated, this is propagated through the network. Conventionally layers in a model are considered by depth, however as each union of predicates adds additional information to a belief, this is referred to as the fidelity of a belief. Activated atomic beliefs are combined recursively to activate or partially activate higher fidelity beliefs. Any node which has been activated or partially activated can be considered a relevant belief. At the output layer all relevant beliefs are passed to the activation function.

First the atomic predicates are activated in the input layer:

For $a_i \in A$,

$$A_{input} = \begin{cases} s, & \text{if } a_i \in C \\ 0, & \text{if } a_i \notin C \end{cases} \tag{8}$$

Higher fidelity beliefs are recursively activated:

$$\forall b \in B, \quad \text{Activation}(b) = w_b \sum_{c \in b} \text{Activation}(c) \quad \text{if} \quad b \cap C \neq \emptyset$$

Resulting in a set of Relevant Beliefs that are used to calculate the output node values:

$$R = \{b \in A \cup B \mid \text{Activation}(b) > 0\}$$

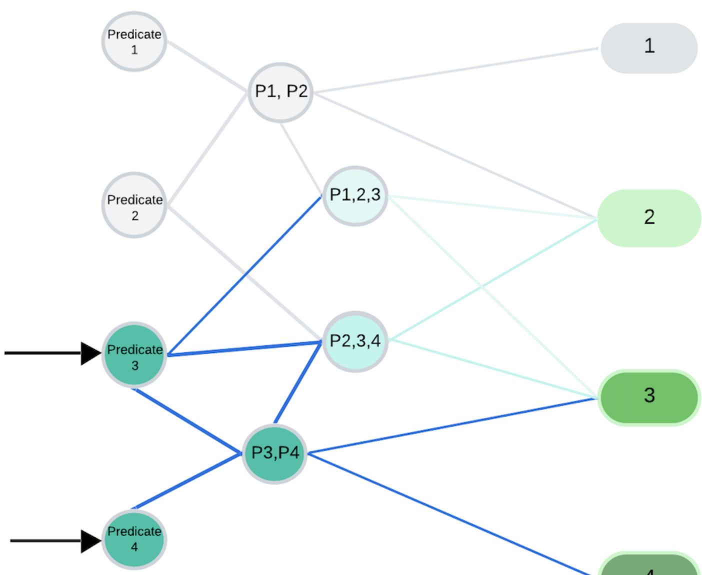

Fig. 4. An example of the dynamic activation based on relevant beliefs, and how this propagates through the model

5.1.7. Model training

Once the model is built, the model can then be trained using conventional back propagation techniques [69]. It uses each instance in the training set, and activates the relevant beliefs by passing a scaled distance value, represented within the context graph, where $i_n \equiv c_x^{ins_m}$, then propagates the output through to the output nodes. The truth value is the target label, which is compared with the output values, such that:

$$\text{loss}(t) = \text{Cross Entropy}\left(\phi_{tn}^{\text{ins}_m}, \max(\phi_t)\right) \tag{9}$$

The loss is then propagated using an optimisation algorithm, such as Adam [27], against the parameters existing within the nodes and edges. As only the relevant nodes were activated, the gradients outside these nodes will be zero, therefore not affected. It is important to note that each instance is assessed against each target in its list individually. The model can hold multiple output targets (represented as multiple layers) but the individual forward pass through the model is assessed against a single target object, as such the loss is taken from ϕ_{tn} and not Φ. The model, as with all learning networks is heavily influenced by the learning rate, to support additional nuances a learning rate scheduler is integrated into the model.

In a divergence from conventional neural networks, belief-net has auxiliary training modes to supplement the conventional supervised learning approach described above. The instance based learning method is in effect an online learning approach, enabling an agent to shift the confidence in different beliefs through the lifecycle of the model, and conduct key tasks such as domain adaption. This is achieved through the retention of the optimiser state through life, meaning that it can be used to conduct back-propagation once with additional instances on a case by case basis. Key to this would be a method for the agent to self-supervise and label its own samples. In this approach the learning rate remains consistent with the learning rate scheduler. Where this is not available, there is also a manual learning mode, which enables the model to train with human labelled samples, using a similar approach. The key separation is that an operator can specify the learning rate. These two functions create the training capabilities which enables an autonomous agent to adapt through life and through domains.

5.1.8. *Output generation*

The generation of an output also holds some key separations from a conventional neural network. As previously mentioned, there is no matrix multiplication conducted as part of the inference process. Whilst this could have a performance impact, this is offset by the overall sparsity of the model, for a given inference there may only be a small proportion of the overall model activated at anyone time. However, the output nodes still need to draw from the precursor nodes to formulate an output. This is done through recursive node activation, in which each node calls back through the network, extracting the outputs from their predecessors and calculates a node output. This function, is called each time inference is run, similarly to a conventional predict function. In integrating this function, the model is able to account for new predicates, beliefs and relationships to be integrated into the model. Uniquely BeliefNet uses an output layer per prediction object, which provides the model with its generalisation performance. Each layer has weighted relationships, and means that new output layers can be integrated into model without having any direct experience of an object and make generalised assessments. The algorithm detailing how outputs are generated is shown in 2

5.1.9. *Ontology*

The model is not dependent upon a fixed ontology, but as it is predicate based, commonality in an ontology between perception and reasoning modules maximises model performance. It is important to note, that the model could accept inputs from sensor modules with new, different or limited predicate sets. For this example we will use the a layered ontology which was derived using the Yamaha CMU dataset [62] and used to re-label the input images to BeliefNet. This enables object class, environmental meta-data and class properties to be analysed by BeliefNet. The use of a dense layered ontology, with referable relationships, can lead to high reasoning performance. The ontology is hierarchical, with classes, subclasses and types, this enables the extrapolation of general knowledge which can be integrated as knowledge into the model. By way of example, if we have a subclass of 'terrain', 'grass' with three types, 'low', 'medium' and 'tall'.

$$\begin{aligned}
&g = \{g_l, g_m, g_t\} \\
&\text{If } \langle g_l, e, \phi_{tn}\rangle \text{ exists and } \langle g_m, e, \phi_{tn}\rangle \text{ exists,} \\
&\text{then } \langle g, e, \phi_{tn}\rangle \text{ is created with a manually specified weighting } w_{\text{manual}}.
\end{aligned} \tag{10}$$

Algorithm 2 Output Generation in BeliefNet

function GENERATEPREDICTION
  for all layers $t \in T$ **do**
    for all nodes n in layer t **do**
      if ϕ_{tn} has predecessors **then**
        for all predecessor nodes p of ϕ_{tn} **do**
          if ϕ_{tn} not computed yet **then**
            $\phi_{tn} \leftarrow$ CALCULATENODEOUTPUT(p)
          else
            $\phi_{tn} \leftarrow$ cached value
          end if
        end for
      end if
    end for
    if t is the prediction object **then**
      $\phi_t \leftarrow$ SOFTMAX(ϕ_t)
      $O \leftarrow$ ARGMAX($n \in t$)
      return O
    end if
  end for
end function

5.1.10. Long vs short term memory

An emergent property of the BeliefNet architecture is the ability to draw parallels between the approach and human memory. The model extracts the useful information from a large number of dense representations (labelled images in this instance) and converts them to a single dynamic model. Initially there is likely to be a large period of growth in the number of nodes and connections, however the growth will be logarithmic in nature. This means that although the system will continue to gain knowledge and strengthen it's beliefs, it will not require exponential amounts of memory. Whilst neural networks are able to collapse the training data into the model structure, BeliefNet has two benefits which make it more analogous to the conversion of short term to long term memory. First, information from various sources of perception can be integrated into a single representation, which can be directly accessed when when required [41]. Secondly, perhaps of greater distinction, is the persistent nature of beliefs within the architecture. Unlike a connectionist approach, the beliefs are generated once, then evolve in a Bayesian manner, this opposes the conventional approach to neural network development, which due to the non-deterministic nature of it's internal representation, requires a full model retrain at the point new information is available. Combined, these components provide the basic framework to provide an autonomous agent an evolving world model.

5.1.11. Explainability

A key feature of the model structure is the inherent traceability through the model to determine the factors which have led to a given prediction. This can be advantageous in highly regulated domains or environments where human machine collaboration may be high. The traceability is a direct by product of avoiding fully connected layers, meaning that an individual belief or input node can be simply and deterministically assessed for its contribution to a given output. The model nodes retain their previous outputs in state, meaning that a critical path to prediction can be traced from each output node through to the input node by recursively presenting the highest n contributors. This has been integrated directly into the model as an explainability function. A representation of this can then be visualised, as shown in 5. Contribution Cn for a given node i to a subsequent node j is defined as the combination of the output ϕ and the intra-node weight w:

$$Cn_{ij} = \phi_i \times w_{ij} \tag{11}$$

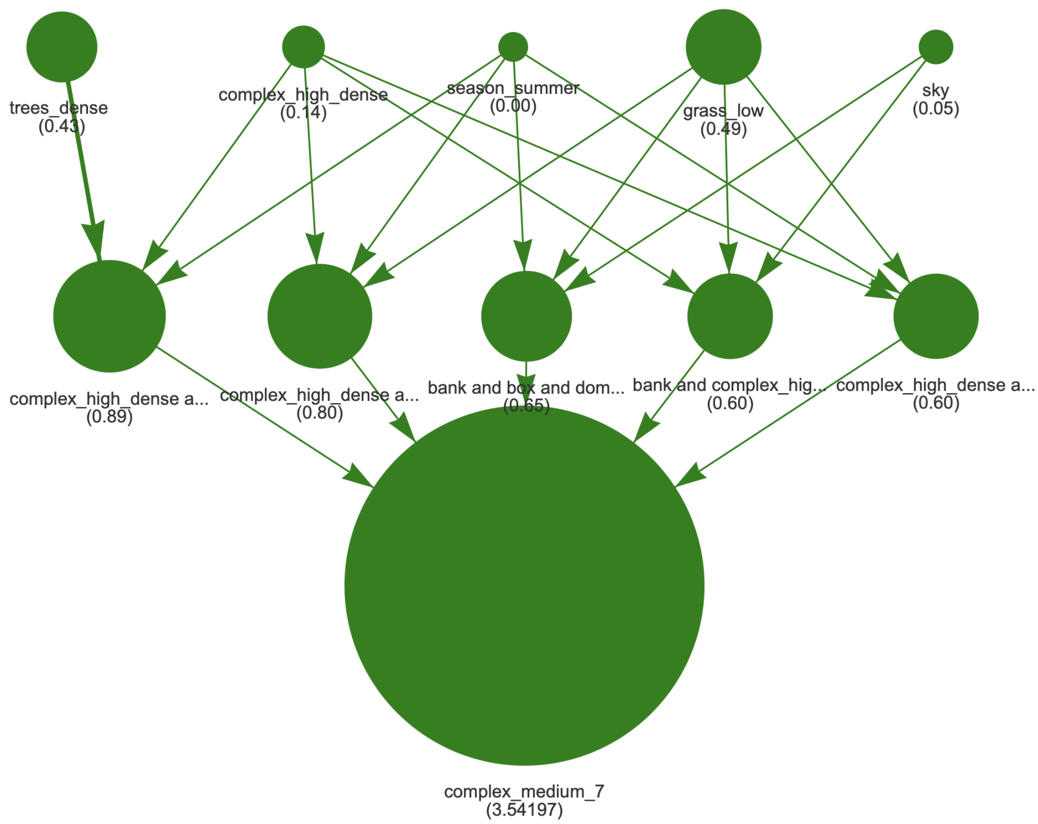

Fig. 5. The models graph explanation function showing the top 5 critical path contributors to the overall output, this is visualised graphically. Contributions are calculated recursively, with each layer showing the contribution to the subsequent node.

6. Experimentation

To test the BeliefNet approach we applied the model to a traversability scenario in which is was presented with a pre-segmented and labelled image, and sought to correctly classify the traversability of specific objects within the image. Within this scenario, we sought to test two factors:

- Terrain classification comparison: how does BeliefNet compare to a static value approach, a conventional graph embedding approach and a random forest classifier.
- Data size comparison: how does BeliefNet compare to graph embedding and random forests as the size of the training data increases.
- Activation function comparison: how does the model adapt with differenct combinations of activation function across the model layers.

A comparative test to an end-to-end model was not conducted, due to a reliance upon both Lidar and imagery for most approaches, and the comparison of a segmented classification and pixel/voxel classification is not a simple translation.

6.1. Approach

The data was formed of the Yamaha CMU dataset, augmented with the ontological based labels [62], and a baseline traversability value (as seen in figure 3) was assigned to each object class. This acted as a baseline, as it accurately represents the current terrain classification approach to traversability assessment, by directly allocating a value to a given class. To build a ground truth dataset, 300 instances were re-labeled with human assessments of the traversability value, enabling humans to extract the image context and make a reasoned assessment on the relative risk associated with each object. This ground truth data is used as the basis for training the BeliefNet. On a physical platform, this approach may still be viable to support manual training enable a human to 'fine-tune' the performance of an agent in a given scenario, or could be replaced directly with an inductive approach in which the platform verifies an assessment based on vehicle dynamics during interaction. The re-labelled samples are then randomly split to provide a training and a test set, with all test metrics being completed by evaluating the test set. This established the framework from which the experimentation was conducted.

The test was targeted at generated responses for the 'grass', 'hardcore', 'soil', 'sand', 'paved' ontology objects, which are the primary traversable objects. Some instances have multiple target objects, meaning that in total there were c.350 training samples. This is a relatively small amount for a traditional complex network, but represents a reasonable amount of varied terrain data that an autonomous system could gather in its initial pre-training for a given domain, similarly it represents the data that could be gathered about a given domain. It enables us to test the ability of the model to adapt to smaller perturbations in the domain and data. The test set was extracted as 20% of the overall training set. At all points in the test this was used to ensure comparability. The random samples were then taken from the training set in increasing increments from 25 samples to the full dataset, and models for each set were trained. Each model was then tested against the test set and the accuracy was judged on the correct categorisation of the risk value against the human adjusted value. This was repeated 15 times and averaged for each model, with a new random test set identified for each iteration. The data holds large variation, due to its size, randomly selecting test data through multiple iteration ensures a variety of complexity, especially zero-shot samples, is tested representatively.

6.2. Metrics

As the overall classification metrics in this instance are risk based and incremental, the performance can also be assessed by assessing the distance in separation between the predicted and actual values. A model that gets its predictions closer to the actual classification performs better than one which is further away. To capture this, we will look at both an absolute classification, but also a fuzzy accuracy which assesses the score as +/- 1 of the absolute.

6.3. Variables

The baseline accuracy is using the default values for an object based upon its ontological class and value, compared against the human edited values, this would be heavily skewed by sampling, so a consistent baseline from the full dataset was taken as 23% absolute accuracy and 43% with fuzzy accuracy.

In addition to the baseline values we tested three additional approaches:

- BeliefNet model as described in this paper.
- A random forest classifier [8], which was chosen to as a comparator due to its reasoning capacity with small datasets, and its ability to explain its results, making it the most similar in output to BeliefNet
- A graph embedding model, which is the closest to a conventional neural network, which uses the GraphSAGE algorithm [19] to create a context graph embedding and then passes the embedding to an XGBoost algorithm, acting as a classification head, to conduct supervised classification [12].

6.4. Outputs

When Belief-Net was trained to predict the outputs of the traversable object classes in the ontology (grass, hardcore, soil, sand, complex, rock), using the full dataset it achieved 47% absolute and 81% fuzzy accuracy, this did not include apriori knowledge. When scaled with the dataset, this performed in as shown in figure 7. This test was repeated with only the grass objects, as these present the greatest proportion of the dataset and are terrain features with the greatest traversability index variation within the class, the results of which are shown in figure 8. The best

Model	Prediction objects	Absolute accuracy	Fuzzy Accuracy
Baseline	All	23%	43%
GraphSAGE + XGBoost	All	33%	52%
Random Forest	Grass	35%	72%
	All	49%	79%
BeliefNet	Grass	35%	75%
	All	47%	81%

Table 1

The summary results using a full dataset over 15 iterations with random test sets for each of the test models.

performing average results for each model are shown in table 2. The comparison between the baseline and BeliefNet against multiple prediction objects can be seen in figure 6, noting that the number of samples is not consistent across object types this is related to the increased variance in some objects over others, for example the low grass distribution is significantly lower than tall grass. In the more challenging object, tall grass, due to the higher variation, BeliefNet outperformed the baseline in both the absolute an fuzzy accuracy.

The graph embedding model failed to learn effective patterns within the data, this is likely due to the additional abstraction generated by the embeddings and the small amount of data for a given prediction, preventing the model from being able to generalise effectively. This resulted in the model returning the same value for instances of a given terrain, and not identifying any factors which would shift the risk. Even training using the full dataset, the model returned an absolute score of 33% and fuzzy score of 52%. The graphSAGE model is the comparator to a conventional neural network, the inability to converge on a solution demonstrates the importance of a neuro-symbolic approach in a complex reasoning task.

The random forest was more successful and was able to make comparable predictions in both actual and fuzzy accuracy, with the full training data achieving 79% fuzzy accuracy, compared with 81% for Belief-Net, as shown in figure 8. Further more random forests present two additional downsides when compared to the BeliefNet model.

The nature of random forests, mean that it is challenging for them to form predictions across multiple classifications and classification objects. As a result, each classification object, i.e. grass_low, required its own model. Whilst this is a standard practice it comes with a number of drawbacks, firstly it prevents generalised concepts being formed across multiple terrain types, in effect reducing the training data available to each model, this will impact domain adaption. Secondly in practice, there will be an i/o cost to loading new models, which could be a bottleneck in situations with more than one traversable object. Given the volume of assessments required in a continuous terrain classification this will have a significant cost. Some instances in the dataset had 5 target objects, meaning 5 separate models would need to be loaded for 1 image. In contrast the BeliefNet is capable of having multiple output layers simultaneously for a single model backbone. This means the model is able to draw generalised concepts rather than terrain specific ones, which provides significant advantages, when the domain ontology adapts. This can be seen in the data, a random forest was trained for each object, meaning that throughout the training it has always seen a representation of the object previously, where as it is possible that BeliefNet makes classifications with no prior knowledge of an object. In all evaluation runs the BeliefNet would make a prediction on atleast one class that was not in its training distribution. This represents a trade-off between accuracy and generalisation, and is demonstrated clearly by the separation between absolute accuracy in all prediction objects. Although this is 3% separation, it is likely that this is the benefit of having a specific model for each class. Whilst this is beneficial, this is out-weighted significantly by the model being able to make predictions on unseen dataclasses, as the BeliefNet demonstrates.

An additional advantage of BeliefNet over random forest relates to the fixed inputs required for a random forest. The input data for the model is a fixed shape array with each item in the array reflecting an possible context object, and the distance from that object. This has two drawbacks, firstly in an ontology such as the one used in this model, with greater than 70 objects, this results in a very sparse set of input data, which can lead to over fitting [67] and may be a contributor to the flat learning profile. Secondly, the fixed nature means that the model cannot adapt to new objects identified within the domain. If a new object was identified, based on a new or adapted sensor classifier, the model would require retraining. In contrast, BeliefNet has a dynamic input length, requiring only the predicates that are sensed to be passed, and it is designed to be extensible, when a new predicate is identified, this can be directly

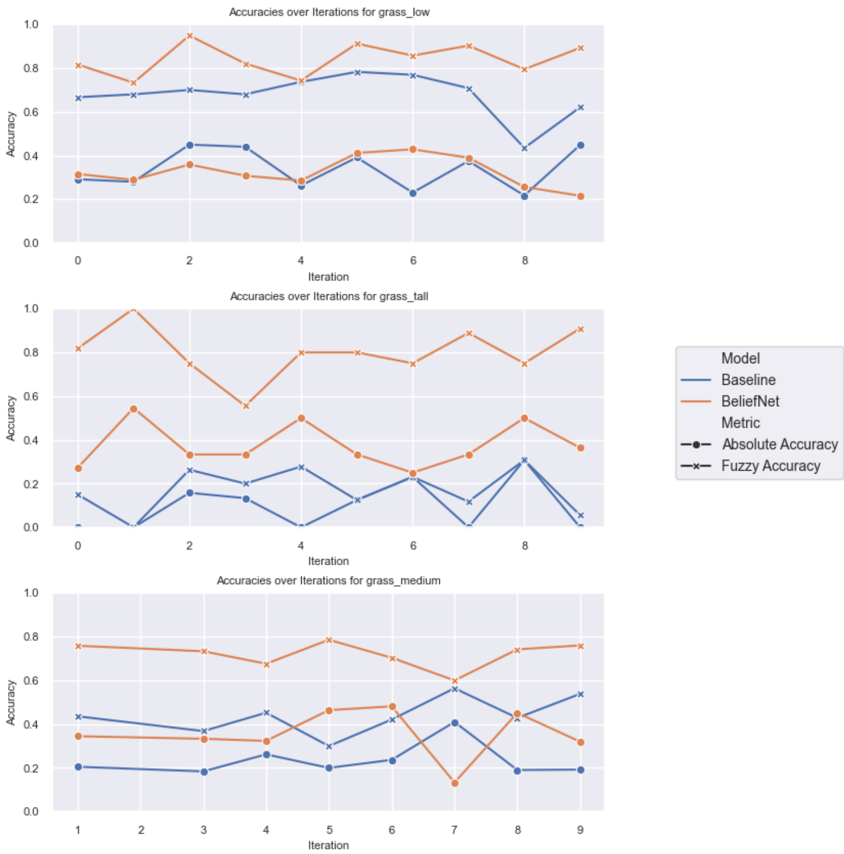

Fig. 6. Experiment comparison of Belief-Net and the baseline absolute and fuzzy accuracy for individual grass prediction objects, the variance in prediction value increases in objects left to right.

integrated into the model. In this case weights are initialised with a default value, but can then be fine-tuned, but in a manner which constrains the adaption only to the relevant predicates, as only they are activated. This prevents having an adverse impact on existing and unrelated concepts. This flexibility and adaptive structure is core to BeliefNets domain generalisation and establishes it as a through-life model, which grows with the agents understanding of the world.

To validate the performance characteristics of the model, we tested the grass sample set using a number of activation functions, in doing so we are able to see how the model adapts over different combinations. Activation functions were assigned to the input layer and belief nodes separately, noting they each held separate behaviours. A number of functions were used:

- Leaky-Rectified Linear Unit.
- Linear activation, in effect the identity of the input.
- Hard-sigmoid.
- Learnable Rectified Linear Unit (ReLU), which was generated with a learnable scalar parameter p such that $\mathrm{ReLU}(x) * P$.

The experiment sought to identify any key variations in the results from the separate activation functions. Each combination was repeated 15 times and the mean results are shown in figure 9, using a consistent learning rate of 0.001 and over 15 epochs of learning. The model performed consistently across the models. The best performing combinations were those which entailed a linear function at the input layer. As the input is a function of distance, this suggests the model benefits from retaining the symbolic information. The learnable activation functions performed

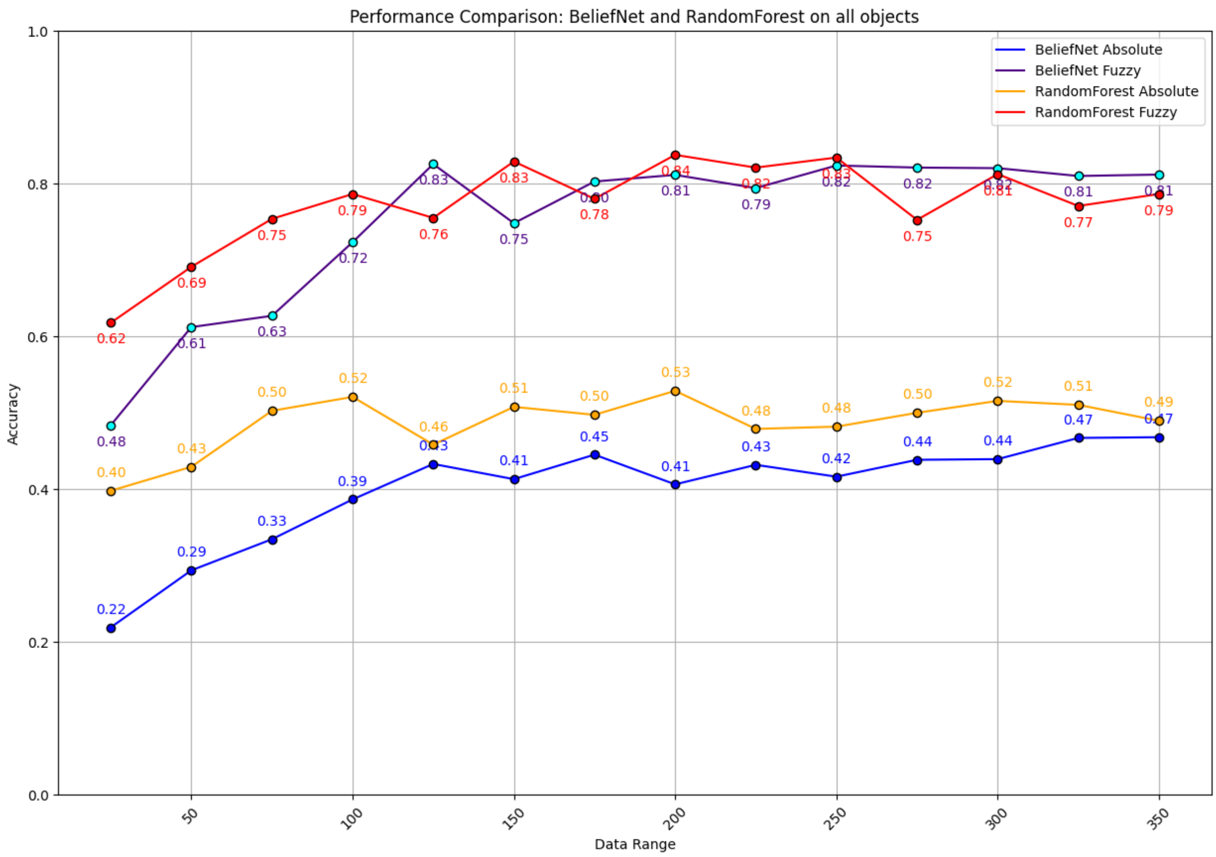

Fig. 7. Experiment comparison of Belief-Net and a Random Forest, with a scaled dataset comparing classification of objects 'grass', 'sand', 'hardcore', 'complex', 'soil' risk classifications

Table 2

Model	Fuzzy Acc Var	Abs Acc Mean	Abs Acc Var	Abs Acc Std Dev	Fuzzy Acc Mean	Fuzzy Acc Std Dev
Linear-HardSigmoid	0.003302	0.353033	0.000431	0.020770	0.801267	0.057464
Learnable-LeakyReLU	0.000445	0.327467	0.000934	0.030568	0.735500	0.021083
LeakyReLU-LeakyReLU	0.007640	0.324967	0.004002	0.063263	0.713000	0.087407
Linear-Learnable	0.001807	0.367900	0.004646	0.068162	0.814167	0.042511
HardSigmoid-Learnable	0.004052	0.368600	0.002058	0.045362	0.751433	0.063652

well, but not significantly outperforming, suggesting there are sufficient model parameters without the requirement to augment.

6.5. Further comparison

As stated above this paper has not contrasted the BeliefNet directly with an end-to-end model in performance, but it presents a number of additional benefits outside of performance, such as explainability, domain adaption as previously highlighted. But it also presents an opportunity to be modality agnostic, meaning sensors and classification

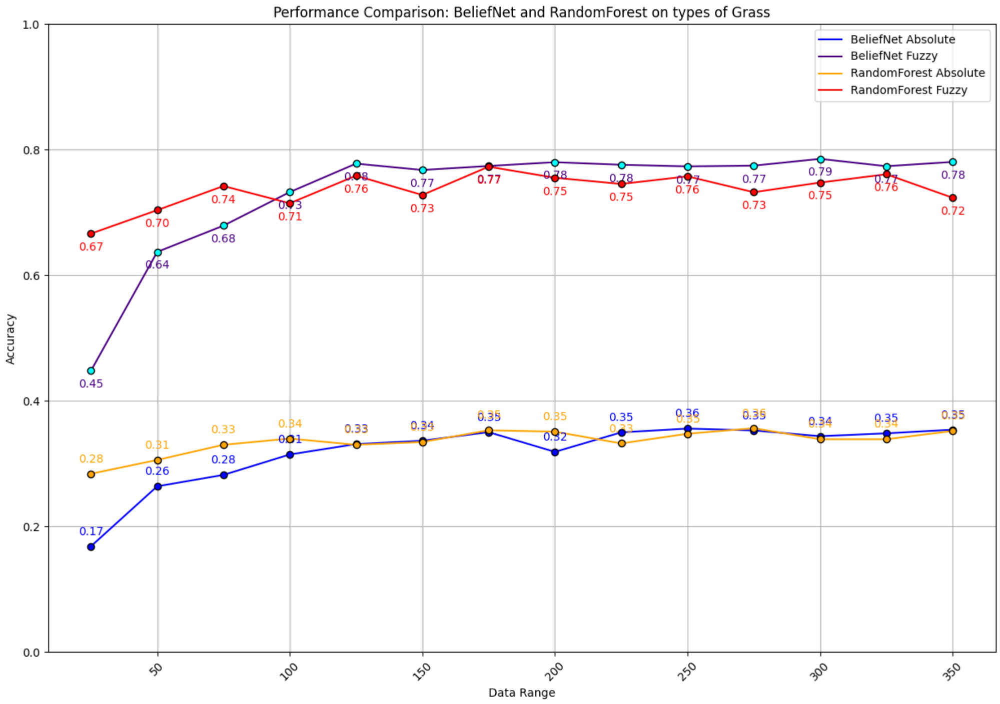

Fig. 8. Experiment comparison of Belief-Net and a Random Forest, with a scaled dataset comparing classification of Grass objects *grasslow*, *grassmedium*, *grasstall* risk classifications

models can be exchanged, and if they adhere to a common ontology can be directly integrated into the reasoning engine. By dislocating the reasoning and sensors, it means agent performance can retain consistency in performance (to a degree) whilst being able to increase sensing performance.

6.6. Further work

This work outlines the potential for BeliefNet in the domain of complex environment traversability, but there are opportunities for further development which could enhance its applicability. Firstly, research into the perception module conducting the initial classification was out of scope for this paper, extending the solution further to include the connection of a single, or multi-modality sensor module would be the next towards platform integration. Secondly, this experiment also out scoped the platforms conversion of exteroceptive and interoceptive sensing outputs into a traversability assessment, thus creating a full learning loop for the agent, this could be another application of BeliefNet. In addition BeliefNet provides a model which is tolerant to an expanding ontology, defining how this might be achieved in a reactive manner before feeding into the network, would be a valuable extension. Another area to be considered is experimentation with the learning rate for manual learning/human intervention, such that learning is effective, without adversely skewing model outputs. Finally this research into the models performance was completed against a single objective function, expanding the research to support multi-objective optimisation would enable additional agency in more complex situations. For example, the ability for BeliefNet to support the risk/time trade-off when assessing tactical route planning.

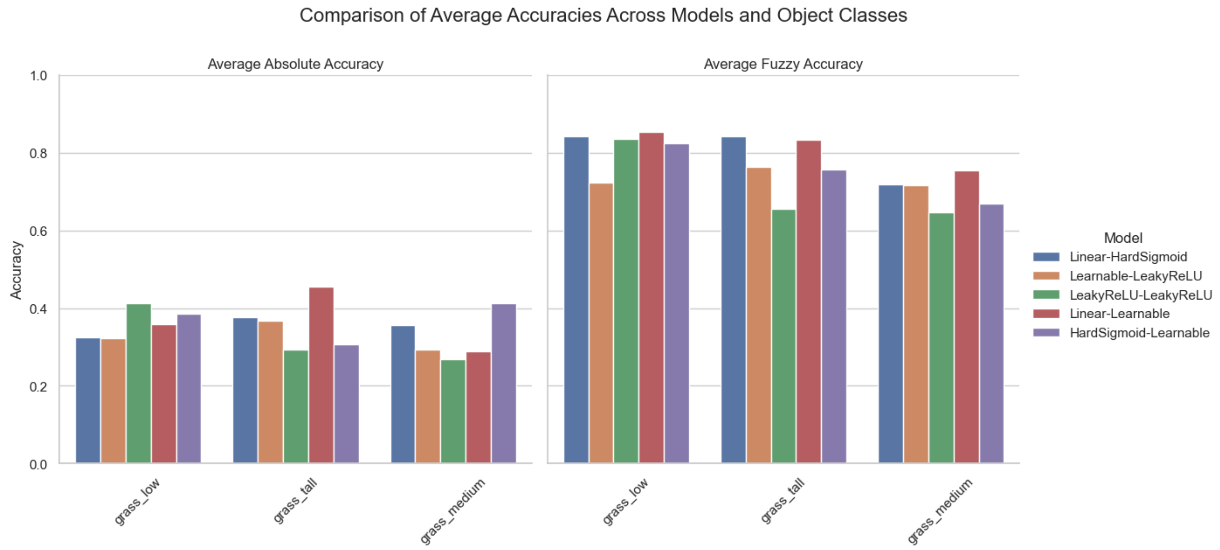

Fig. 9. Experiment comparison of Belief-Net using separate activation functions (input, belief node) when classifying the grass objects(*grasslow*, *grassmedium*, *grasstall*)

7. Summary and conclusions

BeliefNet presents an opportunity to provide a unified reasoning engine to support terrain traversal, in a manner which enables an agent to make an informed decision about risk and traversability. It is inherently extensible meaning it can use what it has learnt within one domain, and adapt this to unknown environments and its inherent explainability means that operators and interpret, understand and impact decision making. This approach significantly increases performance when compared with the static value approach, and enhances the flexibility and explainability when compared to an end-to-end model. This paper demonstrated the application of the BeliefNet model within an autonomous agent traversability reasoning task, however this model structure has potential to applied more widely across similar tasks with high complexity and underlying logic.

If we return to the four questions posed in the introduction of this paper, BeliefNet may not be the complete answer to what is a very complex area, but presents an important step towards autonomous reasoning in complex environments. Unlike an end-to-end model, static values or other ML models, it presents a method of developing a single world model for a single objective. It is able to adapt, learn and crucially evolve the structure of the model, rather not just its weights, enhancing its ability to operate in unseen environments. It can present a traceable outline to the prediction, making it explainable and auditable. Finally this was based purely on the output of a vision model, which presents a very useful opportunity for passive sensing in constrained domains such as Defence. This model presents an opportunity to advance collaboration across human machine teams, enabling operator to integrate knowledge and interact with the training process. This tackles two of the key challenges, it increases the performance of the agent, whilst enabling users to increase their trust through engagement and understanding. BeliefNet presents an advance towards enabling autonomous systems to reasons around complex environments, then learn, adapt and evolve with experience.

References

[1] A. Agarwal, P. Shenoy and a. Mausam, End_to_End Neuro_Symbolic Architecture for Image_To_Image Reasoning Tasks, *arXiv* (2021). http://arxiv.org/pdf/2106.03121.

[2] A. Agha, K. Otsu, B. Morrell, D.D. Fan, R. Thakker, A. Santamaria-Navarro, S.-K. Kim, A. Bouman, X. Lei, J. Edlund, M.F. Ginting, K. Ebadi, M. Anderson, T. Pailevanian, E. Terry, M. Wolf, A. Tagliabue, T.S. Vaquero, M. Palieri, S. Tepsuporn, Y. Chang, A. Kalantari, F. Chavez, B. Lopez, N. Funabiki, G. Miles, T. Touma, A. Buscicchio, J. Tordesillas, N. Alatur, J. Nash, W. Walsh, S. Jung, H. Lee, C. Kanellakis, J. Mayo, S. Harper, M. Kaufmann, A. Dixit, G. Correa, C. Lee, J. Gao, G. Merewether, J. Maldonado-Contreras, G. Salhotra, M.S. Da Silva, B. Ramtoula, Y. Kubo, S. Fakoorian, A. Hatteland, T. Kim, T. Bartlett, A. Stephens, L. Kim, C. Bergh, E. Heiden, T. Lew, A. Cauligi, T. Heywood, A. Kramer, H.A. Leopold, C. Choi, S. Daftry, O. Toupet, I. Wee, A. Thakur, M. Feras, G. Beltrame, G. Nikolakopoulos, D. Shim, L. Carlone and J. Burdick, NeBula: Quest for Robotic Autonomy in Challenging Environments; TEAM CoSTAR at the DARPA Subterranean Challenge, arXiv, 2021, arXiv:2103.11470 [cs]. doi:10.48550/arXiv.2103.11470. http://arxiv.org/abs/2103.11470.

[3] S. Badreddine, A. Garcez, L. Serafini and M. Spranger, Logic Tensor Networks, *Artificial Intelligence* (2022), 103649. doi:10.1016/j.artint.2021.103649. http://arxiv.org/abs/2012.13635.

[4] E. Barnes and J. Hutson, Natural Language Processing and Neurosymbolic AI: The Role of Neural Networks with Knowledge-Guided Symbolic Approaches, 2024. https://www.researchgate.net/profile/James-Hutson-4/publication/378461007_Natural_Language_Processing_and_Neurosymbolic_AI_The_Role_of_Neural_Networks_with_Knowledge-Guided_Symbolic_Approaches/links/65d9f941adc608480ae817b8/Natural-Language-Processing-and-Neurosymbolic-AI-The-Role-of-Neural-Networks-with-Knowledge-Guided-Symbolic-Approaches.pdf.

[5] S. Beycimen, D. Ignatyev and A. Zolotas, A comprehensive survey of unmanned ground vehicle terrain traversability for unstructured environments and sensor technology insights, *Engineering Science and Technology, an International Journal* **47** (2023), 101457. doi:10.1016/j.jestch.2023.101457. https://www.sciencedirect.com/science/article/pii/S2215098623001350.

[6] S.F. Bhat, R. Birkl, D. Wofk, P. Wonka and M. Müller, Zoedepth: Zero-shot transfer by combining relative and metric depth, *arXiv preprint arXiv:2302.12288* (2023).

[7] B.P. Bhuyan, A. Ramdane-Cherif, R. Tomar and T.P. Singh, Neuro-symbolic artificial intelligence: a survey, *Neural Computing and Applications* **36**(21) (2024), 12809–12844. doi:10.1007/s00521-024-09960-z.

[8] L. Breiman, Random Forests, *Machine Learning* **45**(1) (2001), 5–32. doi:10.1023/A:1010933404324.

[9] J. Bridle, Training Stochastic Model Recognition Algorithms as Networks can Lead to Maximum Mutual Information Estimation of Parameters, in: *Advances in Neural Information Processing Systems*, Vol. 2, Morgan-Kaufmann, 1989. https://proceedings.neurips.cc/paper/1989/hash/0336dcbab05b9d5ad24f4333c7658a0e-Abstract.html.

[10] F. Castaño, G. Beruvides, R.E. Haber and A. Artuñedo, Obstacle Recognition Based on Machine Learning for On-Chip LiDAR Sensors in a Cyber-Physical System, *Sensors* **17**(9) (2017), 2109, Number: 9 Publisher: Multidisciplinary Digital Publishing Institute. doi:10.3390/s17092109. https://www.mdpi.com/1424-8220/17/9/2109.

[11] R.O. Chavez-Garcia, J. Guzzi, L.M. Gambardella and A. Giusti, Learning Ground Traversability From Simulations, *IEEE Robotics and Automation Letters* **3**(3) (2018), 1695–1702. doi:10.1109/LRA.2018.2801794. https://www.researchgate.net/profile/R-Chavez-Garcia/publication/319897323_Learning_Ground_Traversability_From_Simulations/links/5ad9b0c3a6fdcc293586e005/Learning-Ground-Traversability-From-Simulations.pdf.

[12] T. Chen and C. Guestrin, XGBoost: A Scalable Tree Boosting System, arXiv, 2016, arXiv:1603.02754. doi:10.48550/arXiv.1603.02754. http://arxiv.org/abs/1603.02754.

[13] L.N. DeLong, R.F. Mir, M. Whyte, Z. Ji and J.D. Fleuriot, Neurosymbolic AI for Reasoning on Graph Structures: A Survey, *arXiv preprint arXiv:2302.07200* (2023). https://arxiv.org/pdf/2302.07200.

[14] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit and N. Houlsby, An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale, arXiv, 2021, arXiv:2010.11929 [cs]. doi:10.48550/arXiv.2010.11929. http://arxiv.org/abs/2010.11929.

[15] D.D. Fan, K. Otsu, Y. Kubo, A. Dixit, J. Burdick and A.-A. Agha-Mohammadi, STEP: Stochastic Traversability Evaluation and Planning for Safe Off-road Navigation, *Robotics: Science and Systems* (2021).

[16] P. Filitchkin and K. Byl, Feature-based terrain classification for LittleDog, in: *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, IEEE, Vilamoura-Algarve, Portugal, 2012, pp. 1387–1392. ISBN 978-1-4673-1736-8 978-1-4673-1737-5 978-1-4673-1735-1. doi:10.1109/IROS.2012.6386042. http://ieeexplore.ieee.org/document/6386042/.

[17] L. Fritz, H.A. Hamersma and T.R. Botha, Off-road terrain classification, *Journal of Terramechanics* **106** (2023), 1–11. doi:10.1016/j.jterra.2022.11.003. https://www.sciencedirect.com/science/article/pii/S0022489822000751.

[18] P. Groumpos, New Scientific Field for Modelling Complex Dynamical Systems: The Cybernetics Artificial Intelligence (2024). NewScientificFieldforModellingComplexDynamicalSystems:TheCyberneticsArtificialIntelligence.

[19] J. Hawthorne, Inductive Logic, in: *The Stanford Encyclopedia of Philosophy*, Spring 2021 edn, E.N. Zalta, ed., Metaphysics Research Lab, Stanford University, 2021. https://plato.stanford.edu/archives/spr2021/entries/logic-inductive/.

[20] M. Himmelsbach, F.v. Hundelshausen and H.-J. Wuensche, Fast segmentation of 3D point clouds for ground vehicles, in: *2010 IEEE Intelligent Vehicles Symposium*, 2010, pp. 560–565, ISSN: 1931-0587. doi:10.1109/IVS.2010.5548059.

[21] T.S. Hodgdon, A.J. Fuentes, J.L. Olivier, B.G. Quinn and S.A. Shoop, Automated terrain classification for vehicle mobility in offroad conditions, Report, Cold Regions Research and Engineering Laboratory (U.S.), 2021, Accepted: 2021-04-08T16:13:35Z Artwork Medium: PDF Interview Medium: PDF Publication Title: This Digital Resource was created in Microsoft Word and Adobe Acrobat. doi:10.21079/11681/40219. https://erdc-library.erdc.dren.mil/jspui/handle/11681/40219.

[22] M. Ingram, R. Moreton, B. Gancz and F. Pollick, Calibrating Trust Toward an Autonomous Image Classifier · Volume 2, Issue 1, 2021. https://tmb.apaopen.org/pub/calibrating-trust/release/2.

[23] H. Inotsume and T. Kubota, Terrain traversability prediction for off-road vehicles based on multi-source transfer learning, *ROBOMECH Journal* **9**(1) (2022), 6. doi:10.1186/s40648-021-00215-3.

[24] G. Jocher, Ultralytics YOLOv5, 2020. doi:10.5281/zenodo.3908559. https://github.com/ultralytics/yolov5.

[25] Kahneman and Daniel, Thinking, Fast and Slow, 2017. http://dspace.vnbrims.org:13000/jspui/bitstream/123456789/2224/1/Daniel-Kahneman-Thinking-Fast-and-Slow-.pdf.

[26] H. Kautz, The Third AI Summer: AAAI Robert S. Engelmore Memorial Lecture, *AI Magazine* **43**(1) (2022), 105–125, Number: 1. doi:10.1002/aaai.12036. https://ojs.aaai.org/index.php/aimagazine/article/view/19122.

[27] D.P. Kingma and J. Ba, Adam: A Method for Stochastic Optimization, 2014. https://arxiv.org/abs/1412.6980v9.

[28] H. Kokel, The 6 Types of Neuro-Symbolic Systems, 2020, Section: posts. https://harshakokel.com/posts/neurosymbolic-systems/.

[29] L.C. Lamb, A. Garcez, M. Gori, M. Prates, P. Avelar and M. Vardi, Graph Neural Networks Meet Neural-Symbolic Computing: A Survey and Perspective, arXiv, 2021, arXiv:2003.00330 [cs]. doi:10.48550/arXiv.2003.00330. http://arxiv.org/abs/2003.00330.

[30] J. Larson and M. Trivedi, Lidar based off-road negative obstacle detection and analysis, in: *2011 14th International IEEE Conference on Intelligent Transportation Systems (ITSC)*, 2011, pp. 192–197, ISSN: 2153-0017. doi:10.1109/ITSC.2011.6083105.

[31] Y. LeCun,
A Path Towards Autonomous Machine Intelligence Version 0.9.2, 2022-06-27.

[32] Y. Lecun, U. Muller, J. Ben, E. Cosatto and B. Flepp, Off-Road Obstacle Avoidance through End-to-End Learning., 2005.

[33] L. Luo, G. Zhang, H. Xu, Y. Yang, C. Fang and Q. Li, INSIGHT: End-to-End Neuro-Symbolic Visual Reinforcement Learning with Language Explanations,
arXiv preprint arXiv:2403.12451. https://arxiv.org/pdf/2403.12451.

[34] R. Manduchi, A. Castano, A. Talukder and L. Matthies, Obstacle Detection and Terrain Classification for Autonomous Off-Road Navigation, *Autonomous Robots* **18**(1) (2005), 81–102. doi:10.1023/B:AURO.0000047286.62481.1d.

[35] J. Mao, C. Gan, P. Kohli, J.B. Tenenbaum and J. Wu, The Neuro-Symbolic Concept Learner: Interpreting Scenes, Words, and Sentences From Natural Supervision, arXiv, 2019, arXiv:1904.12584 [cs]. doi:10.48550/arXiv.1904.12584. http://arxiv.org/abs/1904.12584.

[36] D. Maturana, P.-W. Chou, M. Uenoyama and S. Scherer, Real-Time Semantic Mapping for Autonomous Off-Road Navigation, in: *Field and Service Robotics*, Vol. 5, M. Hutter and R. Siegwart, eds, Springer International Publishing, Cham, 2018, pp. 335–350. ISBN 978-3-319-67360-8 978-3-319-67361-5. http://link.springer.com/10.1007/978-3-319-67361-5_22.

[37] H. Mercier and D. Sperber, Why do humans reason? Arguments for an argumentative theory, *Behavioral and Brain Sciences* **34**(2) (2011), 57–74. doi:10.1017/S0140525X10000968. https://www.cambridge.org/core/product/identifier/S0140525X10000968/type/journal_article.

[38] A. Mileo, Towards a Neuro-Symbolic Cycle for Human-centered Explainability | Neurosymbolic Artificial Intelligence,
Neurosymbolic Artificial Intelligence. https://neurosymbolic-ai-journal.com/paper/towards-neuro-symbolic-cycle-human-centered-explainability.

[39] A. Mileo, Towards a neuro-symbolic cycle for human-centered explainability, *Neurosymbolic Artificial Intelligence* (2024), 1–13. doi:10.3233/NAI-240740.

[40] B. Millidge, A. Seth and C.L. Buckley, Predictive Coding: a Theoretical and Experimental Review, arXiv, 2022, arXiv:2107.12979 [cs, q-bio]. http://arxiv.org/abs/2107.12979.

[41] D.E. Nee, M.G. Berman, K.S. Moore and J. Jonides, Neuroscientific Evidence About the Distinction Between Short- and Long-Term Memory, *Current Directions in Psychological Science* **17**(2) (2008), 102–106. doi:10.1111/j.1467-8721.2008.00557.x.

[42] L. Newman, Descartes' Epistemology, in: *The Stanford Encyclopedia of Philosophy*, Winter 2023 edn, E.N. Zalta and U. Nodelman, eds, Metaphysics Research Lab, Stanford University, 2023. https://plato.stanford.edu/archives/win2023/entries/descartes-epistemology/.

[43] P.K. Panigrahi and S.K. Bisoy, Localization strategies for autonomous mobile robots: A review, *Journal of King Saud University - Computer and Information Sciences* **34**(8, Part B) (2022), 6019–6039. doi:10.1016/j.jksuci.2021.02.015. https://www.sciencedirect.com/science/article/pii/S1319157821000550.

[44] C. Pearson, Autonomous solutions to get critical supplies to combat zones, 2022. https://des.mod.uk/des-investigating-autonomous-solutions-to-reduce-risk-to-life-of-getting-critical-supplies-to-combat-zone/.

[45] A. Pickering, The Cybernetic Brain: Sketches of Another Future, 2023. https://www.docdroid.net/file/download/spEaFSo/pickering-cybernetic-brain-pdf.pdf.

[46] M.-C. Popescu, V.E. Balas, L. Perescu-Popescu and N. Mastorakis, Multilayer perceptron and neural networks, *WSEAS Trans. Cir. and Sys.* **8**(7) (2009), 579–588.

[47] M. Qu and J. Tang, Probabilistic Logic Neural Networks for Reasoning, arXiv, 2019, arXiv:1906.08495 [cs, stat]. doi:10.48550/arXiv.1906.08495. http://arxiv.org/abs/1906.08495.

[48] J. Redmon and A. Farhadi, YOLO9000: Better, Faster, Stronger, in: *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE, Honolulu, HI, 2017, pp. 6517–6525. ISBN 978-1-5386-0457-1. doi:10.1109/CVPR.2017.690. http://ieeexplore.ieee.org/document/8100173/.

[49] R. Riegel, A. Gray, F. Luus, N. Khan, N. Makondo, I.Y. Akhalwaya, H. Qian, R. Fagin, F. Barahona, U. Sharma, S. Ikbal, H. Karanam, S. Neelam, A. Likhyani and S. Srivastava, Logical Neural Networks, 2020. https://arxiv.org/abs/2006.13155v1.

[50] J. Seo, S. Sim and I. Shim, Learning Off-Road Terrain Traversability with Self-Supervisions Only, *IEEE Robotics and Automation Letters* **8**(8) (2023), 4617–4624, arXiv:2305.18896 [cs]. doi:10.1109/LRA.2023.3284356. http://arxiv.org/abs/2305.18896.

[51] H. Seraji and A. Howard, Behavior-based robot navigation on challenging terrain: A fuzzy logic approach, *IEEE Transactions on Robotics and Automation* **18**(3) (2002), 308–321. doi:10.1109/TRA.2002.1019461. http://ieeexplore.ieee.org/document/1019461/.

[52] C. Sevastopoulos, K.M. Oikonomou and S. Konstantopoulos, Improving Traversability Estimation Through Autonomous Robot Experimentation, in: *Computer Vision Systems*, D. Tzovaras, D. Giakoumis, M. Vincze and A. Argyros, eds, Springer International Publishing, Cham, 2019, pp. 175–184. ISBN 978-3-030-34995-0.

[53] A. Sheth, K. Roy and M. Gaur, Neurosymbolic AI – Why, What, and How, *arXiv preprint arXiv:2305.00813* (2023). http://arxiv.org/pdf/2305.00813.
[54] D. Silver, J.A. Bagnell and A. Stentz, Learning from Demonstration for Autonomous Navigation in Complex Unstructured Terrain, *The International Journal of Robotics Research* **29**(12) (2010), 1565–1592, Publisher: SAGE Publications. doi:10.1177/0278364910369715.
[55] S. Siva, M. Wigness and J.G.R.H. Zhang, Robot Adaptation to Unstructured Terrains by Joint Representation and Apprenticeship Learning, in: *Robotics: Science and Systems XV*, Vol. 15, 2019. ISBN 978-0-9923747-5-4. p30.pdf.
[56] W. Stacy, J.V. Cohn, A. Geyer and T. Wheeler, A Cognition-Based Control System for Autonomous Robots, *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* **54**(19) (2010), 1473–1477. doi:10.1177/154193121005401925.
[57] Y. Su, T. Wang, S. Shao, C. Yao and Z. Wang, GR-LOAM: LiDAR-based sensor fusion SLAM for ground robots on complex terrain, *Robotics and Autonomous Systems* **140** (2021), 103759. doi:10.1016/j.robot.2021.103759. https://www.sciencedirect.com/science/article/pii/S0921889021000440.
[58] UK MOD, Defence Artificial Intelligence Strategy, 2022. https://www.gov.uk/government/publications.
[59] A. Valada, J. Vertens, A. Dhall and W. Burgard, AdapNet: Adaptive semantic segmentation in adverse environmental conditions, in: *2017 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, Singapore, Singapore, 2017, pp. 4644–4651. ISBN 978-1-5090-4633-1. doi:10.1109/ICRA.2017.7989540. http://ieeexplore.ieee.org/document/7989540/.
[60] M. Visca, S. Kuutti, R. Powell, Y. Gao and S. Fallah, Deep Learning Traversability Estimator for Mobile Robots in Unstructured Environments, arXiv, 2021, arXiv:2105.10937 [cs]. http://arxiv.org/abs/2105.10937.
[61] W.J. von Eschenbach, Transparency and the Black Box Problem: Why We Do Not Trust AI, *Philosophy & Technology* **34**(4) (2021), 1607–1622. doi:10.1007/s13347-021-00477-0.
[62] W. Wang, Yamaha-CMU Off-Road Dataset - AirLab, 2021. https://theairlab.org/yamaha-offroad-dataset/.
[63] G.N. Wilson and A. Ramirez-Serrano, Terrain Roughness Identification for High-Speed UGVs, *Journal of Automation and Control Research* (2014). doi:10.11159/jacr.2014.002. http://jacr.avestia.com/2014/002.html.
[64] S. Xie, R. Song, Y. Zhao, X. Huang, Y. Li and W. Zhang, Circular Accessible Depth: A Robust Traversability Representation for UGV Navigation, arXiv, 2022, arXiv:2212.13676 [cs]. http://arxiv.org/abs/2212.13676.
[65] J. Yang, Y.Z. Ang, Z. Guo, K. Zhou, W. Zhang and Z. Liu, Panoptic Scene Graph Generation, *arXiv preprint arXiv:2207.11247* (2022). doi:10.48550/arXiv.2207.11247. https://arxiv.org/pdf/2207.11247.
[66] S.H. Young, Robot autonomy in complex environments, in: *Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications III*, Vol. 11746, SPIE, 2021, p. 1174602. doi:10.1117/12.2591478.
[67] J. Zhang and W. Lu, Sparse data machine learning for battery health estimation and optimal design incorporating material characteristics, *Applied Energy* **307** (2022), 118165. doi:10.1016/j.apenergy.2021.118165. https://www.sciencedirect.com/science/article/pii/S0306261921014380.
[68] Justified True Belief - an overview | ScienceDirect Topics. https://www.sciencedirect.com/topics/psychology/justified-true-belief.
[69] Learning from Demonstration for Autonomous Navigation in Complex Unstructured Terrain. doi:10.1177/0278364910369715.