Optimisation of Disaster Resource Distribution with Risk Uncertainty Using Fuzzy Theory

Diego Perez-Palacin
Linnaeus University
Växjö, Sweden
Karlsruhe Institute of Technology
Karlsruhe, Germany
diego.perez@lnu.se

Kenneth Johnson
Auckland University of Technology
Auckland, New Zealand
kenneth.johnson@aut.ac.nz

Vincenzo Grassi
Università di Roma Tor Vergata
Rome, Italy
vincenzo.grassi@uniroma2.it

Raffaela Mirandola
Karlsruhe Institute of Technology
Karlsruhe, Germany
raffaela.mirandola@kit.edu

Abstract

Disasters are unexpected events that can cause widespread disruption of an urban region. Disaster managers are tasked with distributing lifesaving resources to meet evolving demand in the immediate aftermath of a disaster. However, decision-making during the first response is often characterised by a lack of information on current road conditions, and thus, the travel risk posed to teams remains uncertain. This paper presents a decision-making framework underpinned by Markov decision processes to verify provably correct team resource distribution strategies satisfying temporal logic constraints under travel risk vagueness. We implement the framework to express vagueness using fuzzy logic and demonstrate its effectiveness across a range of disaster resourcing scenarios.

CCS Concepts

• **Computing methodologies → Planning under uncertainty**; *Markov decision processes.*

ACM Reference Format:
Diego Perez-Palacin, Kenneth Johnson, Vincenzo Grassi, and Raffaela Mirandola. 2026. Optimisation of Disaster Resource Distribution with Risk Uncertainty Using Fuzzy Theory. In *IEEE/ACM 14th International Conference on Formal Methods in Software Engineering (FormaliSE '26), April 12–13, 2026, Rio de Janeiro, Brazil.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3793656.3793691

1 Introduction

Disasters occur suddenly and are often unexpected. They can cause harm to people, have negative impact on health and cause damage to infrastructure. It is the role of disaster managers to command and coordinate teams of first responders to distribute scarce resources such as food, water and medicine throughout the disaster-affected region. Resourcing operations are supported through the traditional disaster management lifecycle, with four primary stages: preparedness, mitigation, response and recovery [26].

The *resourcing problem* is a critical lifesaving operation to provide resources where needed and involves two steps in different stages of the disaster management lifecycle. First, an initial resource *allocation* is performed during the preparedness stage where resources are allocated satisfying carefully defined requirements specifying amounts of resources likely needed at each location within the region. Second, a resource *distribution* occurs during the response stage. In this step, disaster managers command and coordinate first response teams to move resources within the region to where they are in demand. Their decisions advise teams to avoid damaged routes, which may cause them to fail, whenever possible.

Formal verification techniques provide correctness guarantees for both steps of the resourcing problem. Finding a suitable initial allocation in the first step boils down to the constraint solving problem $\alpha \models K$, whereby resourcing demands are translated into the predicate K and an allocation is modelled by the resource assignment function α, following [13]. The high level of uncertainty surrounding the disaster's impact on the region makes resource distribution decisions much more challenging for disaster managers, especially in the initial stages when only limited information is available [5, 25, 27, 29]. Markov Decision Processes (MDPs) [10] are non-deterministic mathematical models for computing optimal action decisions within a system exhibiting stochastic behaviours through probabilistic model checking of temporal formulae codifying updated resource demands. Verification results obtained from model checkers [11, 20] are optimally *safe* resource distribution strategies for teams which avoid severely disaster impacted routes [12–14].

However, a drawback to typical Markov models is that they assume availability of precise probability values to measure the risk of traversing a damage impacted route. These values may not always be available to disaster managers and even if they are, they can quickly become outdated in a rapidly evolving disaster response scenario. Our paper addresses this problem by using fuzzy logic theory [17] to handle ambiguity and vagueness of information inherent in disaster risk. Leveraging established theoretical results on MDP [10] and on Fuzzy Linear Programming problems (LPP) [28], we specify resource distribution as a Fuzzy LPP to optimise

FormaliSE '26, Rio de Janeiro, Brazil

ACM ISBN 979-8-4007-2478-7/26/04
https://doi.org/10.1145/3793656.3793691

team resource distribution actions within the region, where risk is quantified as fuzzy numbers. Our new approach offers realistic support for decision-makers and provides a range of possible optimal strategies to allow further analysis to be conducted when more information becomes available over the course of the disaster response.

Summarising, our main contributions are:

- bespoke transformations from the resourcing problem with vague information to MDP with fuzzy probabilities and to Fuzzy LPP
- a method for resolving Fuzzy LLPs whose constraints contain fuzzy numbers representing probabilities to obtain optimally safe resource distribution strategies
- a comparison of the strategies obtained with vague information against those obtained via MDP verification with certain information, and
- experimental results of the Fuzzy LLP across a range of resource distribution scenarios with an open-source implementation of our approach.

The rest of the paper is structured as follows. Section 2 develops an algebraic specification of the resourcing problem. Team decision making is formalised as a Markov Decision Process in Section 3 and illustrated with a case study. Section 4 translates the Fuzzy Linear Programming Problem (LPP) from the MDP. Section 5 shows experimental results. Section 6 summarizes the related works. Section 7 discusses complexity, generalizability and tuning of the approach. Lastly, Section 8 give concluding remarks and outlines future research.

2 Resourcing Problem Specification

During the preparedness phase of the disaster management lifecycle, managers specify quantities of emergency *resources* to initially allocate to specific locations within the region. However, during disaster response and recovery, resources must be moved throughout the region according to additional demands for resources. This is challenging due to rapidly evolving changes in resource requirements as the disaster scenario progresses. In this section, we follow [14, 15, 21] to give an algebraic specification of the *resourcing problem*, comprising two primary tasks for disaster managers: initial allocation of resources according to a set of requirements, and resource distribution within the region, as demand evolves over time. Let $R = \{R_1, \ldots, R_m\}$ be the set of m sorts of resources such as food, water, medicine and building supplies.

To illustrate the resourcing problem, Figure 1 introduces a running example showing a portion of the Auckland Central Business District (CBD) overlaid with the graph $G = (V, E)$ such that $V = \{a, \ldots, h\}$ are vertices representing locations where resources can be placed, interconnected with edges from the set E. The graph G is bidirectional such that for each edge we have $u \longrightarrow v \in E \iff v \longrightarrow u \in E$. Edges are labelled with traffic light symbols `R:Red`, `Y:Yellow`, and `G:Green` from the set I representing disaster impact severity. Distinguishing among three levels of risk in a route is a reasonable trade-off in the first moments after a disaster, when information quantity and quality is limited, between not discerning any risk in the routes and a complete characterization of each route. First response teams prefer to travel on safe (`Green`) edges rather than higher-risk (`Red` or `Yellow`) ones. Symbols in I can be modified and new ones added to express more complex road conditions. The total function $i : E \longrightarrow I$ assigns each edge in $e \in E$ an impact symbol $i(e) \in I$.

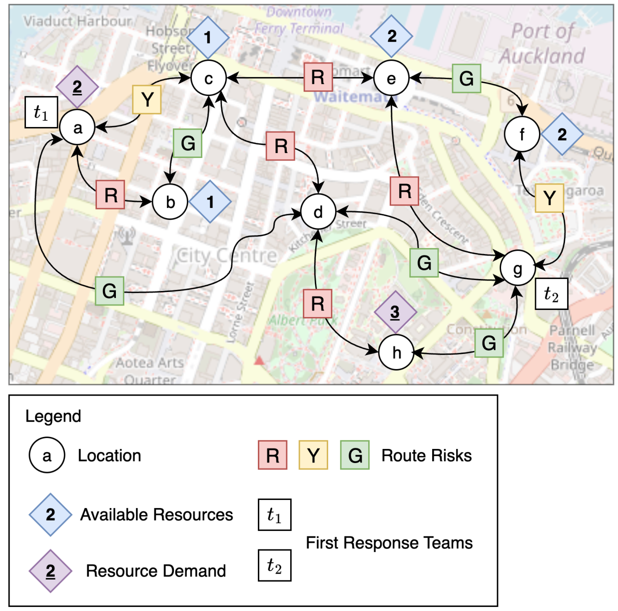

Figure 1: Auckland CBD OpenStreetMap overlaid with graph G

2.1 Initial Resource Allocation

Prior to a disaster, managers place resources according to carefully designed requirements which specify, at least initially, the amount of each sort of resource that will likely be needed at each location, should a disaster occur. The goal of the initial resource allocation is to work out the correct amounts of resources per location while adhering to particular resourcing requirements that prescribe how many of a particular resource sort should be placed at a location.

Let X be the set of variables of sort $\mathbb{N}$ and define the natural number sorted variable $\mathtt{xv}$ for each vertex v in V. The function $\alpha : X \longrightarrow \mathbb{N}^m$ is an *assignment function* such that the equation $\alpha(\mathtt{xv}) = (a_1, \ldots, a_m)$ means location v has the amount $a_i \in \mathbb{N}$ of resource sort r_i. To simplify our discussion, we set $m = 1$. Next, we fix a number $N \in \mathbb{N}$ to define the finite-value assignment function $\alpha : X \longrightarrow \{0, \ldots, N\}$, such that $0 \leq \alpha(x) \leq N$ for all $x \in X$.

Previous works [13, 15] developed a formal methodology based on predicate logic over the natural numbers to automatically calculate an initial resource assignment function α satisfying resourcing requirements $Q_1, \ldots, Q_n$. For example, disaster managers might require

- Q_1: Locations `b` and `c` have the same amount of resources
- Q_2: Location `e` has at least two resources, and
- Q_3: Location `f` has more resources than `c` and `e` combined.

Each resourcing requirement can be translated into a predicate over variables in X such that

- K_1: $\mathsf{xb} = \mathsf{xc}$
- K_2: $\mathsf{xe} \geq 2$
- K_3: $\mathsf{xf} > \mathsf{xc} + \mathsf{xe}$

Predicates are built from variables in X using arithmetic operations and defined inductively using relational symbols $>, \geq, =, \leq$ and $<$ in the standard way. There are implicit requirements on variables to ensure their values are non-negative, yielding a constraint $K \equiv \bigwedge_{i=1}^{n} K_i$.

The initial resource allocation task is formalised into the computational task of finding an assignment function α satisfying K. In symbols, $\alpha \models K$. Figure 1 shows an assignment α satisfying $K_1 \wedge K_2 \wedge K_3$ such that $\alpha(c) = \alpha(b) = 1$, $\alpha(e) = 2$ and $\alpha(f) = 2$, setting $\alpha(\mathsf{vx}) = 0$ for all other variables.

2.2 Resource Distribution

In the aftermath of a disaster, first responders must move resources according to updated demands. Let $T = \{t_1, \ldots, t_k\}$ be the set of all teams distributing resources within the region given by the graph G. For example, Figure 1 shows first response teams t_1 and t_2 stationed at locations a and g respectively. The two teams are responsible for moving resources by traversing the edges of G to where they are needed within the region, represented by underlined values $\underline{2}$ and $\underline{3}$ at locations a and h. To capture the current *resourcing state* within the region we specify the tuple

$$\sigma = (G, \rho, \alpha) \tag{1}$$

comprising G, the graph of the region, $\rho : T \rightarrow V$, a function assigning teams in T to vertices in V, and α the assignment function. When $\alpha \models K$ we say that σ satisfies K.

Let S be the set of all resourcing states (1) of the disaster-affected region. We suppose that new resource demands within G are reflected in an update to K that forms new constraints K'. From our running example, we have new constraints $K' \equiv xa = 2 \wedge xh = 3$. The first responders t_1 and t_2 must move appropriate amount of resources via the edges of G to where they are needed. To this end, we define the operation $move_i : S \times V^2 \times \mathbb{N} \rightarrow S$ such that the equation $move_i(\sigma, u, v, k) = \sigma'$ means that the new state σ' is obtained from team t_i moving k resources from $u \in V$ to $v \in V$. We have the following

Definition 2.1 (Operation Validity). The operation $move_i(\sigma, u, v, k)$ is *valid* on state σ when the following conditions are satisfied:

(1) $\rho(t_i) = u$, team t_i is at location u
(2) there exists an edge $u \rightarrow v$ in E, and
(3) $\alpha(u) \geq k$, there are at least k resources at location u.

If σ satisfies Conditions 1 to 3 then the $move_i$ operation is valid, and updates the state $\sigma' = (G, \rho', \alpha')$ according to

$$\rho'(t) = \begin{cases} v & \text{if } t_i = t \\ \rho(t) & \text{otherwise} \end{cases} \tag{2}$$

which means that team t_i's location is updated from u to v after the move has completed. No other team location is modified. The resourcing assignment is updated by setting

$$\alpha'(w) = \begin{cases} \alpha(u) - k & \text{if } w = u \\ \alpha(v) + k & \text{if } w = v \\ \alpha(w) & \text{otherwise} \end{cases} \tag{3}$$

Team	**Edge**	**Action**
Team t_1	$a \rightarrow b$	$move_1(a, b, 0)$
	$b \rightarrow c$	$move_1(b, c, 1)$
	$c \rightarrow a$	$move_1(c, a, 2)$
Team t_2	$g \rightarrow e$	$move_2(g, e, 0)$
	$e \rightarrow f$	$move_2(e, f, 2)$
	$f \rightarrow g$	$move_2(f, g, 3)$
	$g \rightarrow h$	$move_2(g, h, 3)$

Table 1: Correct Strategy satisfying K'

where k is the amount of resources removed from location u and placed at location v. No other resourcing is modified. Similarly, there is no change to the graph G. However, if any of these three conditions are unsatisfied, then there is no change to the state and $\sigma' = \sigma$. The operation $move_i(\sigma, u, v, 0)$ allows team t_i to move from u to v without transporting resources.

We shall omit reference to the state and write $move_i(u, v, k)$ as an *action* to be applied to any valid state. We form the set $Act = \{move_i(u, v, k)\}$ of all actions for teams $t_i \in T$, vertices $u, v \in V$ and $0 \leq k \leq N$. An action $a \in Act$ is valid in a given state if the three conditions in Definition 2.1 are satisfied. Let $A(\sigma) \subseteq Act$ be the set of all actions valid in state $\sigma \in S$.

2.3 Strategy Correctness

Given an initial state σ, we define a resource distribution strategy as a finite sequence π of actions from the set Act. The sequence π is *correct* if, and only if, it comprises valid actions and yields a state σ' in the set of *target states* $Tar(K') \subseteq S$ such that $\sigma' \in Tar(K') \iff \alpha' \models K'$.

For example, we suppose that the target is for the teams to move resources across the region in a way that satisfies

$$K' \equiv (\mathsf{xa} = 2) \wedge (\mathsf{xh} = 3). \tag{4}$$

Table 1 shows one possible strategy for the teams, comprising seven $move_i$ actions. Starting from the initial state σ represented in Figure 1, Team t_1 moves $a \rightarrow b \rightarrow c \rightarrow a$ delivering two resources to a, satisfying constraint ($\mathsf{xa} = 2$). Team 2 subsequently moves $g \rightarrow e$ to collect two resources, and then $e \rightarrow f$ to collect an additional resource from f. Finally, t_2 moves $f \rightarrow g \rightarrow h$, delivering three resources to satisfy constraint ($\mathsf{xh} = 3$). The resulting state σ' satisfies the resource distribution target

$$\alpha' \models (\mathsf{xa} = 2) \wedge (\mathsf{xh} = 3) \tag{5}$$

and we write $\sigma' \in Tar(K')$.

3 Uncertainty in Resource Distribution

While the strategy listed in Table 1 yields a state satisfying K', the teams moved through routes labeled Red, indicating increased risk of failure of the team to achieve its target. For example, actions $move_1(a, b, 0)$ moving Team t_1 from $a \rightarrow b$ and $move_2(g, e, 0)$ moving Team t_2 from $g \rightarrow e$ are reasonable since they offer shorter travel distances, but they do so at increased risk.

To reason mathematically with risk, work in [14] supposed disaster managers assign precise probability values to the symbols in

I to quantify the disaster's impact on the routes. For example, specific probability values are assigned to symbols in I by the function $p : I \longrightarrow [0, 1]$ defined by cases for $i \in I$

$$p(i) = \begin{cases} 0.5 & \text{if } i = \texttt{Red} \\ 0.85 & \text{if } i = \texttt{Yellow} \\ 0.99 & \text{if } i = \texttt{Green} \end{cases} \tag{6}$$

These values measure the probability of $move_i(u, v, k)$ succeeding, with team t_i successfully traversing edge $u \longrightarrow v$.

When an action of Team t_i fails, the team becomes unavailable and the k resources remain at the source u. To model this mathematically, we extend the set of locations V by adding the $\perp$ element to form $V^{\perp} = V \cup \{\perp\}$ and define $\rho : T \longrightarrow V^{\perp}$ as $\rho'(t) = \perp$ if $t_i = t$ and $\rho(t)$ otherwise. When $\rho(t_i) = \perp$, t_i is unavailable.

A drawback to this approach, however, is obtaining these probability values since, after a disaster, there is no clear knowledge of the concrete probability of success when traversing a route. Hence, assigning a fixed probability of success to the risk of Red, Yellow, and Green routes incurs an overconfidence in the current knowledge.

3.1 Fuzzy Sets and Fuzzy Numbers

In this work, we address overconfidence by allowing vagueness and uncertainty in the risk values of each route using fuzzy set theory. *Fuzzy sets* [17] are able to represent vagueness and uncertainty in a system. Let X be a set and $\mathcal{F}(X) = \{A | A : X \longrightarrow [0, 1]\}$ the set of all functions from X to interval [0,1]. A fuzzy set defined over the universe of discourse X is characterized by the *membership function* $A \in \mathcal{F}(X)$, where, given $x \in X$, $0 \leq A(x) \leq 1$ represents the *membership degree* of x to A. Classical sets can be seen as a special case of fuzzy sets, where $A : X \longrightarrow \{0, 1\}$. A fuzzy set $A \in \mathcal{F}(X)$ defines its *support* as $supp(A) = \{x \in X \mid A(x) > 0\}$, and its *core* as $core(A) = \{x \in X \mid A(x) = 1\}$. A *normal* fuzzy set A holds that $core(A) \neq \emptyset$. The α-cut of a fuzzy set A is a crisp set defined as ${}^{\alpha}A = \{x \in X \mid A(x) \geq \alpha\}$.

Fuzzy numbers are able to represent vagueness in the value of a number. A fuzzy number is a fuzzy set A defined on the set of real numbers $\mathbb{R}$ and that holds the following properties [17]:

- A must be a normal fuzzy set
- ${}^{\alpha}A$ must be a closed interval $[{}^{\alpha}A_l, {}^{\alpha}A_u]$, $\forall \alpha \in (0, 1]$
- the support of A must be bounded

Since probability values are bounded in [0,1], a fuzzy number B that represents a vague probability value must also hold that $supp(B) = [s_l, s_u] \subseteq [0, 1]$. Note that the particular case of a crisp real probability r corresponds to the case of a fuzzy probability B where ${}^{\alpha}B = \{r\}$ $\forall \alpha \in (0, 1]$ and $s_l = s_u = r$. Formally, the set of fuzzy numbers between 0 and 1 is defined as $F_{[0,1]} = \{B \in \mathcal{F}(\mathbb{R}) \mid B \text{ is a fuzzy number and } supp(B) = [a, d], 0 \leq a \leq d \leq 1\}$. Figure 2 illustrates this through an example, where the probability value is 0.5 for the Red routes. The left chart depicts that the probability of succeeding in traversing a Red route is 0.5 with complete certainty. In turn, the fuzzy part indicates that 0.5 remains the most plausible value for Red, but it could be any value between 0.4 and 0.6, with a linearly increasing level of plausibility from the extremes towards 0.5. This forms a triangular fuzzy number.

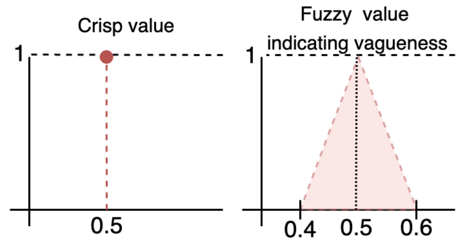

Figure 2: Example crisp and fuzzy value for Red

We therefore redefine the function p to return fuzzy numbers as $p : I \longrightarrow F_{[0,1]}$, and reuse its value to quantify the uncertainty of an action as $p^{Act} : Act \longrightarrow F_{[0,1]}$ to assign a vague probability to each action, defined by $p^{Act}(move_i(u, v, k)) = p(i(u \longrightarrow v))$.

We use in the following p^{Act} to define a transition probability over the states in S to support decision-making in disaster response scenarios.

3.2 Team Markov Decision Process

To reason with team routing decision-making under uncertainty, we define

Definition 3.1. A *Team Markov Decision Process* (Team MDP) is a tuple $M = (S, Act, P)$ where S is the finite set of states, Act the finite set of actions of the form $move_i(u, v, k)$, and P is the probability transition function $P : S \times Act \times S \longrightarrow F_{[0,1]}$, which returns the vaguely defined probability of reaching state σ' from σ by action a.

The transition probability function is defined as:

$$P(\sigma, a, \sigma') = \begin{cases} p^{Act}(a) & \text{if } \sigma' = \sigma'_{SUCC} \\ \boldsymbol{1} - p^{Act}(a) & \text{if } \sigma' = \sigma'_{FAIL} \\ 0 & \text{otherwise} \end{cases} \tag{7}$$

where $\sigma'_{SUCC} = (G, \rho', \alpha')$ corresponds to a successful movement with ρ' and α' defined according to Equations (2) and (3) with Team t_i moving k resources from u to v, where $\sigma'_{FAIL} = (G, \rho', \alpha)$ corresponds to a failed movement of the team, and where $\boldsymbol{1}$ is the fuzzy number that corresponds to the crisp value 1; i.e., the element $\boldsymbol{1} \in F_{[0,1]}$ is defined as $\boldsymbol{1}(1) = 1$ and $\boldsymbol{1}(x) = 0,\ 0 \leq x < 1$. Regarding the subtraction $\boldsymbol{1} - p^{Act}(a)$ with $p^{Act}(a) \in F_{[0,1]}$, we calculate it by applying Zadeh's extension principle. We can easily see that, when $p^{Act}(a) \in F_{[0,1]}$, also $\boldsymbol{1} - p^{Act}(a) \in F_{[0,1]}$, with membership function given by: $(\boldsymbol{1} - p^{Act}(a))(z) = p^{Act}(a)(1 - z), 0 \leq z \leq 1$.

It is worth discussing the requirement in crisp Markov Decision Processes that, for any given state and action, the sum of the probabilities of the outgoing transitions equals one. The restriction introduced in other works on the fuzzification of stochastic systems [2], in order to enable correct modeling and analysis, is that there must exist $cp_{\sigma,a,\sigma'} \in core(P(\sigma, a, \sigma'))$ such that $\sum_{\sigma'} cp_{\sigma,a,\sigma'} = 1$. In the Team MDP with vague probabilities, we have $\sum_{\sigma'} P(\sigma, a, \sigma') = p^{Act}(a) + \boldsymbol{1} - p^{Act}(a) = \boldsymbol{1}$. By the description above of $\boldsymbol{1}$ and of $\boldsymbol{1} - p^{Act}(a)$, we can see that the Team MDP always satisfies such a restriction since $core(\boldsymbol{1}) = 1$.

3.3 Verifying Safe Resource Distribution Guarantees

A *strategy* is a function $\pi : S \longrightarrow A$ that determines the action in $\pi(\sigma) \in A(\sigma)$ to perform in a given state σ [10]. Probabilistic model checkers [11, 20] compute strategies from PCTL reachability properties `Pmax=?[F(`K'`)]` which comprise team actions which maximise the probability of reaching a target state in $Tar(K')$ and successfully deliver resources where they are needed; e.g. such that the updates resource assignment satisfies $\alpha' \models K'$. The strategy resolves non-deterministic choice of resource movement actions in each state such that $|A(\sigma)| = 1$ for each state in S, inducing a Discrete-Time Markov Chain (DTMC) M^{π}. We define a path in M^{π} as a sequence $\bar{\sigma} = \sigma_1\sigma_2\ldots$ of states in S such that $P(\sigma_i, \pi(\sigma_i), \sigma_{i+1}) > 0$ and let *Paths* be the set of all paths.

Our goal is to obtain an optimal strategy π to reach a state σ' that satisfies an updated resourcing demand K'.

Team	Edge	Action
Team t_1	$a \rightarrow c$	$move_1(a, c, 0)$
	$g \rightarrow f$	$move_2(g, f, 0)$
	$f \rightarrow e$	$move_2(f, e, 0)$
	$e \rightarrow f$	$move_2(e, f, 2)$
Team t_2	$f \rightarrow g$	$move_2(f, g, 4)$
	$g \rightarrow h$	$move_2(g, h, 3)$
	$h \rightarrow g$	$move_2(h, g, 0)$
	$g \rightarrow d$	$move_2(g, d, 1)$
	$c \rightarrow a$	$move_1(c, a, 1)$
Team t_1	$a \rightarrow d$	$move_1(a, d, 0)$
	$d \rightarrow a$	$move_1(d, a, 1)$

Table 2: Optimal PRISM Strategy Satisfying K'

3.4 Resource Distribution Strategies

Using the Team MDP definition, we form an MDP M of the Auckland CBD, as shown in Figure 1 there are two teams, and so we set $T = \{t_1, t_2\}$, with an initial resource assignment with $\alpha(b) = \alpha(c) = 1$ and $\alpha(e) = \alpha(f) = 2$. The target is to reach a state σ' satisfying Equation (4). To obtain a resource distribution strategy which maximises the probability of $\sigma' \models K'$ in M, we formulate the PCTL formula

```
Pmax=?[F(xa=2 & xh=3)]                (8)
```

using the `Finally` temporal operator. Applying a recent version of the PRISM [20] probabilistic model checker on the Team MDP M of the Auckland CBD running example in Figure 1 and using the crisp probabilities in Equation (6), the `Pmax=?` operation yields the maximum probability value `0.7906937` of reaching a state such that `(xa=2 & xh=3)` is true, determining the safest strategy π, shown in Table 2. This is a path through the synthesised DTMC M^{π} corresponding to the scenario that none of the actions attempted by teams t_1 and t_2 fail and adds five more team actions compared to the strategy in Table 1 so that teams avoid high risk edges (e.g. $a \rightarrow b$ and $e \rightarrow f$) that are labelled `Red`.

The synthesized DTMC has two absorbing states: one for mission success and the other for mission failure. Each transient state in the DTMC has exactly two outgoing transitions: one for the successful move and the other for the failure of the move.

4 Solving Fuzzy Team MDPs

This section describes the process of analysing the fuzzy MDP to find optimal strategies and to evaluate the strategies to determine the expected probability of success. Among the different methods proposed in the literature to analyse an MDP to find the optimal strategies [10], we selected a method that encodes the Team MDP given in Definition 3.1 as a Linear Programming Problem (LPP) for maximizing the expected probability of reaching any of the target states in $Tar(K')$. We opted for this method as it enjoys a well-established theoretical foundation within fuzzy logic. The LPP encodes an equation for each existing transition between states in the MDP. The set of states that originate transitions to other states, S_{Trans}, contains those ones which are not a target state in $Tar(K')$ and in which at least one of the teams can move because it has not failed yet; i.e, not in the set $S_{FAIL} = \{\sigma((V^{\perp}, E), \rho, \alpha) \mid \rho(p_i) = \perp \forall i\}$. Therefore, $S_{Trans} = S \setminus (Tar(K') \cup S_{FAIL})$.

We encode the Team MDP as an LPP as follows:

Definition 4.1 (Fuzzy Team Linear Programming Problem)**.** Find a z satisfying the objective function

$$\text{minimise } z = \sum_{i=1}^{n} x_i$$

subject to:

$$x_i \geq \sum_{s' \in S \setminus Tar(K')} P(s, a, s') \cdot x_{s'} + \sum_{t \in Tar(K')} P(s, a, t) - \epsilon \quad (9)$$

$\forall s \in S_{Trans}$, $a \in Act$, where $n = |S_{Trans}|$ and $0 < \epsilon$ is a small value $\epsilon << \min_{i \in Act} p^{Act}(i)$ to favour shorter travels with similar probability of success by slightly penalizing each proposed movement.

This definition transforms the transition structure of a Team MDP and the objective of reaching a target state $t \in Tar(K')$ into a linear optimisation problem [10]. However, our problem differs from typical LPPs in that the values of $P(s, a, s')$ and $P(s, a, t)$ in the constraints part in Definition 4.1 do not contain crisp probabilities but fuzzy numbers that represent the uncertainty in routes. The inequalities in Definition 4.1 become an *almost inequality.*

We solve the Fuzzy Team LPP in Definition 4.1 using Tanaka *et al.* method in [28], by defining expressions

$$\begin{aligned}\tilde{Y}_i =& x_0\epsilon + x_i \\ &- \sum_{s' \in S \setminus Tar(K')} P(s, a, s') \cdot x_{s'} - \sum_{t \in Tar(K')} P(s, a, t) \gtrsim 0 \end{aligned} \quad (10)$$

where $\gtrsim 0$ translates to *almost positive* and $x_0 = 1$. The $\tilde{Y}$ forms the typical matrix representation of the LPP constraints $\tilde{Y} = \tilde{A}x \gtrsim 0$, with $\tilde{A}[a_{ij}]$, $a_{i0} = \epsilon$, and only two other elements in each row having a fuzzy value different from 0: the element that corresponds to the destination state if the team movement succeeds and the element that corresponds to the destination state if the team movement fails. In consequence, despite the potentially huge dimension of $\tilde{A}$, it is a sparse matrix and the implementation of the approach greatly benefits from the use of sparse matrices.

The fuzzy numbers in the method in [28] are restricted to the fuzzy numbers defined with two parameters a, b, and the membership function is:

$$A(x) = \begin{cases} 1 - \frac{|a-x|}{b}, \ a-b \le x \le a+b \\ 0, \ \text{otherwise} \end{cases}$$

which correspond to triangular fuzzy numbers. Since we are using fuzzy numbers to represent vague probability values B, we extend the definition of their membership function to express that its vague value is bounded by the interval [0,1]; i.e., $supp(B) = [s_l, s_u] \subseteq [0,1]$, to:

$$B(x) = \begin{cases} 1 - \frac{|a-x|}{b}, \ max(a-b,0) \le x \le min(a+b,1) \\ 0, \ \text{otherwise} \end{cases} \tag{11}$$

where $0 \le a \le 1$ and $b > 0$. For convenience in the analysis, the 0-cut is defined as in [3, 18], in the form ${}^0B = cl\{x \in X | B(x) > 0\}$, where cl denotes the closure of the set. For the type of numbers consider in Equation (11), that definition is equivalent to ${}^0B = [max(a-b,0), min(a+b,1)]$. As illustrative cases of B values, the three extremes are: the case $a = 1$ and $b = 0$ represents the guaranteed success, the case $a = 0$ and $b = 0$ represents the guaranteed failure, and the case $b = \infty$ represents the complete uncertainty about the value since it makes $core(B) = [0,1]$.

To optimise the *almost positive* inequations in Equation (10), we use the algorithm in [28] that iteratively finds the largest possible $h \in [0,1]$ where each fuzzy number in Equation (10) is converted to a crisp value $a - h \cdot b$. Since the fuzzy numbers used in [28] are based on a triangular function, $a - h \cdot b$ corresponds to the lower endpoint of the α-cut where $\alpha = 1 - h$.

In our case, the fuzzy numbers represent vague probability values of successfully moving resources from one location to another, and the resourcing problem does not state any constraint on the minimum expected success probability that can be accepted for a strategy. These characteristics have two implications for the algorithm used. First, the crisp value that corresponds to the lower endpoint of the α-cut, where $\alpha = 1 - h$, cannot be directly instantiated as $a - h \cdot b$ because they must be bounded by [0,1]. They are converted to the crisp value $max(a - h \cdot b, 0)$. Second, a solution to the LPP can still exist up to $h = 1$. The larger the value of h in the interval $[0,1]$, the more pessimistic the analysis is carried out on the routes success probabilities, which results in a lower probability of overall success in the resource distribution strategy, yet a solution exists. Therefore, we save all the strategies that are identified as the optimal strategy for some h. Additionally, we calculate optimal strategies for the upper bound of the α-cut, with $\alpha = 1 - h$, because we are interested in gathering all the possible knowledge given the vague definition of route success probability and not only in a pessimistic analysis. For that, we also compute the optimal strategies instantiating each fuzzy number as $min(a + h \cdot b, 1)$.

Algorithm 1 shows the structure of the instructions for finding the success probability, including vagueness and optimal strategies for the resourcing problem. First, the MDP is created according to the resourcing problem characteristics, and the LPP is encoded from the MDP (lines 1-2). After that, the iterative process for different values of h begins (line 5). In each iteration, the lower endpoint of the α-cut ($\alpha = 1 - h$) is used to instantiate the LPP with

Algorithm 1 Algorithm for finding success probability and optimal strategies in the vague resourcing problem

Require: Parameters map G, constraints K', teams ρ, vague success probabilities p^{Act}, and initial resource distribution α

1: $teamMDP \leftarrow$ createMDP$(K', \alpha, G, \rho, p^{Act})$ *// Definition 3.1*
2: $teamLPP \leftarrow$ encodeLPP$(teamMDP, K', G)$ *// Definition 4.1*
3: $result \leftarrow emptyFuzzyNumber$
4: $optStrat \leftarrow \emptyset$ *// Set of optimal strategies found*
5: **for** $h = 0$ to 1 step Δh **do**
6: $\quad crispLPP \leftarrow$ instantiateLowerEndpointLPP$(teamLPP, h)$
7: $\quad$ buildIncrementalResult($crispLPP$, h, $result$, $optStrat$, $teamMDP$)
8: $\quad crispLPP \leftarrow$ instantiateUpperEndpointLPP$(teamLPP, h)$
9: $\quad$ buildIncrementalResult($crispLPP$, h, $result$, $optStrat$, $teamMDP$)
10: **end for**
11: **Return** result, optStrat

12: **procedure** buildIncrementalResult(crispLPP, h, result, optStrat, teamMDP)
13: $\pi \leftarrow solveLPP(crispLPP)$ *// Find optimal strategy*
14: $optStrat \leftarrow optStrat \cup \pi$
15: $M^{\pi} \leftarrow$ inducedDTMC$(\pi, teamMDP)$
16: $successP \leftarrow$ solveDTMC(M^{π})
17: $result \leftarrow setMembership(result, successP, 1 - h)$
// added membership degree is $result(successP) = 1 - h$
18: **end procedure**

crisp numbers, which is used for solving and finding the optimal strategy for that endpoint value and the success probability of the optimal strategy (lines 6-7). The same operations are done for each upper endpoint (lines 8-9). After the iterations, the *result* stores the fuzzy number that corresponds to the probability of success in the resourcing problem, and *optStrat* is the set of the optimal strategies found. An open-source implementation in MATLAB for the resourcing problem experimented with is available at https://github.com/eres-lnu/fuzzy-resourcing repository.

5 Evaluation

This section experiments with the Fuzzy Team LPP implementation under a variety of resource distribution scenarios. First, we validate our approach by comparing optimal strategies and probability values obtained by the Fuzzy Team LLP to results obtained from probabilistic model checking the corresponding Team MDP when there is no vagueness in the probability; e.g., when $b = 0$ for all fuzzy numbers.

Next, we conduct a series of experiments on the running example presented in Figure 1, with vagueness in the probability for the impact symbols `Red`, `Yellow`, and `Green` represented with the fuzzy numbers: `Red` = (0.5, 0.1), `Yellow` = (0.85,0.35), and `Green` = (0.99, 0.14), for membership functions characterized by (a, b) with the shape defined in Equation (11). Figure 3 illustrates the membership function for each of these fuzzy numbers in their corresponding colours. We conduct experiments over two resource distribution scenarios:

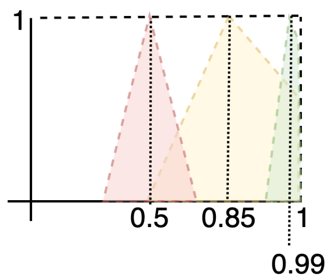

Figure 3: Triangle Diagram Fuzzy numbers representing probability vagueness in impact symbols `Red`, `Yellow`, and `Green`

One Team Scenario comprises a single team, setting $T = \{t_1\}$ which starts at location a, and

Two Teams Scenario comprises two teams, setting $T = \{t_1, t_2\}$ starting in locations a and g, respectively.

Our objective for each experiment is to reach a target state satisfying the resource formulae K' defined in Equation (4). Our experiments calculate optimal strategies for values of h from 0 to 1 in increments $\Delta h = 0.05$.

5.1 LPP Comparison with Probabilistic Model Checking: No Vague Information

Our first experiment applies the Fuzzy Team LLP solution with $b = 0$, when there is no vagueness in probabilities. For the *One Team Scenario*, we solved the Fuzzy Team LLP with the objective of reaching a state in $Tar(K')$ when t_1 starts at location a. We obtained an optimal strategy for coordinating t_1, presented in the leftmost column in Table 3, which succeeds with probability 0.4865. Identical results were obtained from the PRISM probabilistic model checker on the corresponding Team MDP expressed in PRISM's high-level modelling language.

For the *Two Teams Scenario* we compare our Fuzzy Team LPP approach to PRISM probabilistic model checking of the formula (8), on the Team MDP corresponding to two teams t_1 and t_2 starting in locations a and g respectively. Applying the Fuzzy Team LPP with the objective of reaching a state in $Tar(K')$, we obtained an optimal strategy presented in the rightmost column of Table 3, assuming that none of the team actions fail while achieving their objective. The probability of success of reaching a target state is 0.7906, which is identical to the value obtained by PRISM in Section 3.3. However, the optimal strategies differ. PRISM yields a strategy which does not take into account the number of steps performed by the teams, whereas our proposal penalises team movement, and selects an optimal strategy that is logistically more favourable since it is shorter.

5.2 One Team Scenario and Vague Information

In this scenario, we suppose there is only one team $T = \{t_1\}$ available to distribute resources, starting at location a. The resulting probability of success for this scenario is a fuzzy number $P1$ with the shape shown in Figure 4. As expected, the value with the highest membership is $core(P1) = 0.4865$, the same as in our previous experiment in Section 5.1. Looking at the figure to study the

b=0 only t_1	b=0 $t1$ and t_2
$move_1(a, c, 0)$	$move_2(g, f, 0)$
$move_1(c, b, 0)$	$move_2(f, e, 0)$
$move_1(b, c, 1)$	$move_2(e, f, 2)$
$move_1(c, a, 2)$	$move_2(f, g, 4)$
$move_1(a, d, 0)$	$move_2(g, h, 3)$
$move_1(d, g, 0)$	$move_1(a, c, 0)$
$move_1(g, f, 0)$	$move_2(h, g, 0)$
$move_1(f, e, 0)$	$move_2(g, d, 1)$
$move_1(e, f, 1)$	$move_1(c, a, 1)$
$move_1(f, g, 3)$	$move_2(d, a, 1)$
$move_1(g, h, 3)$	

Table 3: Optimal Strategies for no vagueness $b = 0$

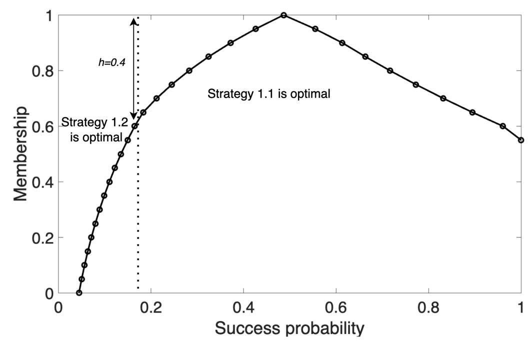

Figure 4: Probability of success for the One Team Scenario.

expectations of a large probability of success, we can see that the membership value of the guaranteed success (success probability 1.0) is $P1(1) \gtrsim 0.55$, while the membership value of success with probability 0.9 increases to $P1(0.9) \simeq 0.65$. In turn, we also look at the figure to study what is most expected to happen. We observe that the 0.5-cut of the fuzzy probability result, which shows that the success value is more expected than not, gives the interval ${}^{0.5}P1 = [0.1220, 1]$. This interval represents a scenario of significant uncertainty for the decision-makers, since it covers almost the entire interval $[0, 1]$. Taking a more restrictive cut to concentrate further on what are the most expected probabilities of success than possible scenarios, the 0.75-cut we obtain ${}^{0.75}P1 = [0.2448, 0.7724]$.

Moreover, in this scenario, there is not a single optimal strategy, but several of them. Consider that, given the level of vagueness on the `Yellow` and `Red`, taking two `Yellow` routes may or may not be more convenient than taking a `Red` route. In a pessimistic situation, where the vaguely defined `Yellow` whose support $supp(Y) = [0.5, 1]$ happens to be close to its lower limit, traversing a `Red` route becomes more convenient than traversing two `Yellow` routes. The Fuzzy Team LPP solution found two different optimal strategies, presented in Table 4.

Strategy 1.1, which is the same strategy as in the no vagueness case, proposes to bring the resources needed in location a from the resources in b and c, and then reach the resources in e and f using a longer yet safer path through d and g to place them

Strategy 1.1	Strategy 1.2
optimal for $0 \leq h < 0.4$ lower endpoints and for $0 \leq h \leq 1$ higher endpoints	optimal for $0.4 \leq h \leq 1$ lower endpoints
$move_1(a, c, 0)$	$move_1(a, c, 0)$
$move_1(c, b, 0)$	$move_1(c, e, 1)$
$move_1(b, c, 1)$	$move_1(e, f, 3)$
$move_1(c, a, 2)$	$move_1(f, g, 5)$
$move_1(a, d, 0)$	$move_1(g, h, 3)$
$move_1(d, g, 0)$	$move_1(h, g, 0)$
$move_1(g, f, 0)$	$move_1(g, d, 2)$
$move_1(f, e, 0)$	$move_1(d, a, 2)$
$move_1(e, f, 1)$	
$move_1(f, g, 3)$	
$move_1(g, h, 3)$	

Table 4: Optimal strategies for one team in function of h

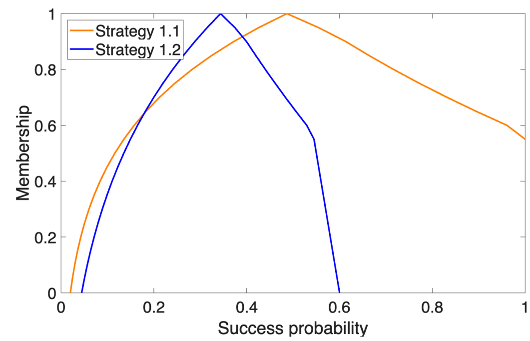

Figure 5: Success probability of optimal strategies for 1-team

in h. This strategy avoids any Red route and is optimal for most possibilities within the vagueness in route risk. Concretely, it is the optimal strategy for the higher endpoints of all h and for the lower endpoints when $h < 0.4$ (i.e, the lower endpoint for ${}^{\alpha}P1$ for $\alpha > 0.6$). However, in the pessimistic situations where the Yellow happens to be towards the most risky possibility of its vague definition, the optimal strategy is different, and it is shown on the right column of Table 4 as Strategy 1.2. It proposes that the Red route between c and e is taken in the second movement to get the two resources in e. Concretely, Strategy 1.2 is optimal for the lower endpoints when $h \geq 0.4$ (i.e., the lower endpoint of ${}^{\alpha}P1$ for $\alpha \leq 0.6$).

Overall, Strategy 1.2 dictates five movements on Green routes, two movements on Yellow routes, and one movement on a Red route, while Strategy 1.1 dictates seven movements on Green routes and four on Yellow routes. The dotted line in Figure 4 separates $P1$ into the probabilities contributed by each of the two strategies that are optimal for at least one situation in the vagueness range. Figure 5 illustrates the success probability when choosing each of the strategies. The expert reader can see that the resulting success fuzzy probability depicted in Figure 4 corresponds to the t-conorm operation of the two fuzzy probabilities in Figure 5 following Gödel logic [17].

5.3 Two Teams Scenario and Vague Information

For this scenario, teams $T = \{t_1, t_2\}$ are on the map at initial locations a and g respectively. When there is vagueness in the route safety, we obtained that the probability of success is a fuzzy number $P2$ with the membership function plotted in Figure 6. We expect a higher success probability than in the previous case because, having two resourcing teams starting in different zones, the chances that risky routes can be avoided increase, as well as the failure of one of the teams traversing a path does not imply the failure of the mission because the other team can take over the remaining part of the mission. Concretely, $core(P2) = 0.7906$, which is aligned with our experiment comparing crisp probability values with probabilistic model checking in Section 5.1.

As in the previous case, we first study the expectations on a high probability of success. The membership value of guaranteed success is $P2 \gtrsim 0.55$, similar to the case with only one team. The reason is that, to offer a guaranteed success, the important aspect is that the paths are safe, not doubling the number of teams. However, the membership value of the rescue success with probability 0.9 increases to $P2(0.9) \gtrsim 0.85$, showing the benefit of having a second team since the previous case was just around 0.65. In turn, the study of what is most expected to happen reveals that the 0.5-cut is ${}^{0.5}P2 = [0.3240, 1]$, and the most expected 0.75-cut is ${}^{0.75}P2 = [0.5282, 0.9594]$, a range of success probabilities more favourable than the previous ${}^{0.75}P1 = [0.2448, 0.7724]$.

For two teams, the probability $P2$ comes from the aggregation of three strategies that are optimal. Table 5 presents the paths of the three strategies in the case where none of the team actions fail. The leftmost column presents Strategy 2.1, which is equivalent to the case without vagueness, where t_2 gathers all four resources from e and f and puts them in g. Afterwards, t_2 delivers from g the three resources required in h, comes back to g, and delivers the remaining resource to a. In turn, t_1 goes to c and brings the resource back to a. Even though the route between c and b is Green and using that resource for the necessities in a would result in fewer total movements (t_2 would not need to reach a, it could remain in h after its delivery), the optimal strategy does not prescribe t_1 reaching b from c to get that resource. The reason is that losing t_1 in that round trip route would force t_2 to do a round trip over the Yellow route a and c to get the resource from c to a. Therefore, the optimal strategy saves effort from t_1 since its successful return from c to a saves t_2 traversing two Yellow routes, at the cost of making t_2 traversing three additional Green routes ($h \rightarrow g \rightarrow d \rightarrow a$). This strategy was found to be the optimal one for the lower endpoints when $h < 0.6$ (i.e., the lower endpoint for ${}^{\alpha}P2$ for $\alpha > 0.4$) and for the higher endpoints when $h < 0.05$.

In situations where the risk of Yellow routes happens to be towards the most risky possibility of its vague definition, the optimal strategy becomes Strategy 2.2, whose fail-free path is shown in the middle column of Table 5. As it happened in the analogous pessimistic case when only t_1 existed, the optimal strategy prefers traversing two Yellow and one Red instead of four Yellow. Having two teams, t_1 moves from a to g by $a \rightarrow c \rightarrow e \rightarrow f \rightarrow g$ gathering all five resources in the way. Once both teams and the resources are in g they split the task. One of the teams delivers two resources to

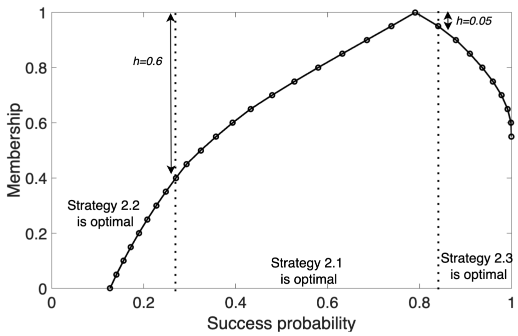

Figure 6: Probability of success for the Two Team Scenario t_1 and t_2.

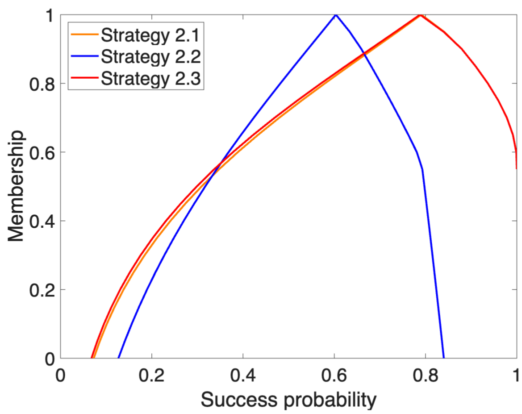

Figure 7: Success probability of optimal strategies for 2-teams

a through $g \to d \to a$, and the other team delivers three resources to h directly from g. At that point, it is indifferent which team delivers to each location. Strategy 2.2 illustrates the scenario where t_1 delivers to h and t_2 delivers to a.

In turn, in optimistic situations where the risk of `Yellow` routes is towards the least risky possibility of its vague definition or when the `Green` routes are completely safe, a new optimal strategy emerges, Strategy 2.3. This strategy prescribes that t_1 gathers resources from both b and c and delivers everything necessary to a. In turn, t_2 gathers the necessary resources from e and f and delivers them to h. Since the `Yellow` and `Green` routes are very safe in this case, t_1 may incur no risk in gathering the resource from b; therefore, it is the preferred option in the optimization, as it favors strategies with fewer movements. The result is that t_2 saves three safe movements at the expense of two safe movements of t_1.

The difference in the optimality of all three strategies in all possible situations, according to the level of vagueness, is shown in Figure 7. The three fuzzy numbers have been created by synthesizing the DTMC that implements each strategy and analysing the probability of reaching the success state using different α-cuts of the vaguely defined success probabilities `Red`, `Yellow`, and `Green`. The figure graphically shows that the fuzzy number for the success probability of Strategy 2.2, $B22$, is, in general, lower than the other two strategies because its peak is more to the left than the success

Strategy 2.1	Strategy 2.2	Strategy 2.3
optimal for $0 \leq h < 0.6$ lower endpoints and for $0 \leq h < 0.0.5$ higher endpoints	optimal for $0.6 \leq h \leq 1$ lower endpoints	optimal for $0.05 \leq h \leq 1$ higher endpoints
$move_2(g, f, 0)$	$move_1(a, c, 0)$	$move_1(a, c, 0)$
$move_1(a, c, 0)$	$move_1(c, e, 1)$	$move_2(g, f, 0)$
$move_2(f, e, 0)$	$move_1(e, f, 3)$	$move_1(c, b, 0)$
$move_2(e, f, 2)$	$move_1(f, g, 5)$	$move_1(b, c, 1)$
$move_2(f, g, 4)$	$move_1(g, h, 3)$	$move_2(f, e, 0)$
$move_2(g, h, 3)$	$move_2(g, d, 2)$	$move_1(c, a, 2)$
$move_2(h, g, 0)$	$move_2(d, a, 2)$	$move_2(e, f, 1)$
$move_2(g, d, 1)$		$move_2(f, g, 3)$
$move_1(c, a, 1)$		$move_2(g, h, 3)$
$move_2(d, a, 1)$		

Table 5: Fail-free paths for the optimal strategies found for two teams in function of h

probability of the other two strategies ($B21$, $B23$) and the membership degree of large values in $B22$ is lower than in $B21$, $B23$ (e.g., $B22(x) < B21(x)$ and $B22(x) < B23(x)$ for $x > 0.8$). Nevertheless, Strategy 2.2 is the preferred strategy in pessimistic situations since the membership degrees of low probabilities of success in $B22$ are lower than in $B21$, $B23$ (e.g., $B22(x) < B21(x)$ and $B22(x) < B23(x)$ for $x < 0.35$). The figure shows that Strategy 2.1 is only slightly better than Strategy 2.3 in unfavourable and normal situations, and both perform equally well in favourable situations ($x > 0.8$). Therefore, Strategy 2.3 is the preferred one only because it comprises fewer team movements.

6 Related Work

In systems characterised by continuously changing conditions, it is essential to adopt modelling approaches that explicitly account for uncertainty [4, 22, 29]. Several studies have applied such approaches to disaster resourcing, including [7, 9, 16, 23]. The resource allocation problem addressed in [16] considers both human and material resources. Post-disaster strategies are calculated using a multi-objective stochastic programming model to minimise unmet resource demand, quantify the probability of a scenario occurring, and determine overall resource needs. Our fuzzy number approach may be used to introduce either vagueness or confidence into these parameters when precise probabilities are unavailable. In [23], allocating and scheduling of rescue teams are computed by a mathematical model to minimize the completion time of relief operations (firefighting, resourcing) and quantifies uncertainty of disaster severity, relief operation times and travel times using fuzzy trapezoidal numbers. Precedence constraints on the ordering of relief operations are also introduced. We plan on extending our approach to consider precedence constraints on operations which fix damaged infrastructure before performing resource distribution operations. A covering tour approach is taken in [7] to distribute resources during the response stage, in which resources can either be taken to where they are demanded or to nearby locations so people may travel to obtain them. An optimisation method is designed based on credibility fuzzy theory, quantifying demand for allocating resources and routing relief vehicles in disaster-affected areas. Currently, our resourcing demands are precise, in the sense that the temporal logic formulae specify exact assignments as in Equation (4). However, we can consider scenarios in which resources may be assigned to any location within a subregion. [9] focuses on disaster relief logistics and optimising placement of relief distribution centres and health centres during the preparedness stage of the disaster lifecycle. In the subsequent response phase, levels of resources to be transferred from supply sources to relief centres and then to where they are needed in the disaster-affected regions are determined by way of multi-objective models with robust stochastic programming. Their models quantify supply resource and cost uncertainties as probability values. While this work focuses on resource distribution, future research will establish a theoretical framework for optimising initial resource allocation during preparedness and for analysing response-stage strategies using predicted resource demand.

Previous research has explored the integration of fuzzy logic and MDPs to analyse decision-making under imprecise information; fuzzifying MDP elements such as states, rewards, and state-action transition probabilities. The fuzziness in associated rewards has been explored, for instance, in [6, 30] by allowing the rewards or costs associated with deciding for a certain action in a given state to be defined as trapezoidal fuzzy numbers [6] and by proposing fuzzy rewards that are risk-sensitive since they are used by risk-averse decision makers [30]. Fuzziness in state transition probabilities in MDPs has been considered in [19], where the fuzzy expected discounted reward for different policies of actions is calculated. MDPs with fuzzy states have been studied in [24], where the best policy to actuate on an M/M/1/N queueing system is calculated, also considering conditional probabilities of fuzzy events upon the occurrence of other fuzzy events. Work in [18] defines a decision model with Markov-type transitions where states, actions, and the discounted total reward are fuzzy. The analysis of DTMCs with fuzzy probabilities is described in [1, 2]. In our work, the DTMCs analysed to calculate the success probability of concrete strategies are absorbing, finite, and have fuzzy transition probabilities.

The fuzzy success probabilities derived for the resourcing problem in this work can be integrated into system-level disaster reasoning, which involves more complex trade-offs among multiple objectives, using fuzzy goal models [8] to capture fuzzy satisfaction.

7 Discussion

In this Section, we discuss the complexity of the proposed method, its generalizability, and a possible model tuning.

Complexity. Regarding the execution time, our experiments indicate that constructing the MDP from the resource allocation problem specification requires ≈8 s, solving each LPP requires ≈33 s, generating each induced DTMC requires ≈0.5 s and solving it ≈0.05 s. All experiments were executed on a MacBook Pro (2021) equipped with an Apple M1 processor and 16 GB of RAM. Regarding space complexity, the resulting 2-teams scenario MDP contains ≈130000 states. The current implementation explicitly generates all states and transitions from the problem specification. The state with the maximum number of outgoing transitions corresponds to the case in which all resources and teams are co-located at a node with the maximum number of neighboring locations (four neighbors in Figure 1). That maximum number is independent of the total number of states and remains constant. Consequently, the space complexity of the MDP is linear in the number of states. Furthermore, the implementation leverages MATLAB's support for sparse matrices when solving the LPPs, ensuring that the memory requirements for both the LPP coefficient matrices and the transition matrices of the induced DTMCs also scale linearly with the number of states.

Generalizability. This work is motivated by the disaster resource distribution problem, offering a concrete setting for studying decision-making under uncertainty with evolving demand and limited environmental information. As discussed in Section 2, although the framework is evaluated in this context, its underlying methodology is domain-independent and applies to a broader class of decision-making problems with similar structures. Disaster response serves as a unifying use case to present the approach end-to-end. The evaluation focuses on a single disaster response scenario, allowing an in-depth analysis of the framework under realistic uncertainty in travel risk and demand. This limits the testing of its generalizability to other contexts involving resource distribution. Nevertheless, the framework's abstract design allows easy adaptation to new scenarios without altering the methodology. A broader assessment across various disasters and conditions remains key future work.

Model tuning. The current model abstracts team failure by returning resources to the source location. In real disaster scenarios, failures may lead to partial or total resource loss. Capturing these effects requires an extended MDP that explicitly models resource loss and degradation in the system state and transitions, increasing complexity. This is a relevant extension planned for future work.

8 Conclusion

We investigated the use of Fuzzy Linear Programming Problems to develop optimally safe strategies for first response teams to allocate lifesaving resources. We chose fuzzy numbers to represent uncertainty because they allowed us to mathematically quantify the vague notions of a disaster's impact on infrastructure, often characterised in the early stages of disaster response. We applied our approach with fuzzy numbers over two resource distribution scenarios. Our results show that when precise probability values are known, we often obtain logically favourable optimal strategies that require less team movement.

We plan to extend the framework to support more expressive constraint specifications, enabling the modeling of realistic resourcing demands over sets of locations. We also plan to incorporate additional operations, such as search and rescue and repair, together with operation precedence constraints. Finally, we will investigate the impact of uncertain initial resource allocations on distribution strategies, requiring the development of enriched Markov models to capture demand uncertainty.

Acknowledgments

This work was partially supported by the Helmholtz Association (HFG) with the KiKIT project and the Grant 46.23 (Engineering Secure Systems) and by the German Research Foundation (DFG) - SFB 1608 - 501798263.

References

[1] James J. Buckley. 2004. *Fuzzy Markov Chains*. Springer Berlin Heidelberg, Berlin, Heidelberg, 35–43. doi:10.1007/978-3-540-36426-9_4

[2] James J. Buckley and Esfandiar Eslami. 2006. Fuzzy Markov Chains: Uncertain Probabilities. https://upcommons.upc.edu/server/api/core/bitstreams/c6e4e981-3a4b-45cf-9f63-6296322aeb16/content

[3] James J. Buckley and Yunxia Qu. 1990. On using α-cuts to evaluate fuzzy equations. *Fuzzy Sets and Systems* 38, 3 (1990), 309–312. doi:10.1016/0165-0114(90)90204-J

[4] Radu Calinescu, Raffaela Mirandola, Diego Perez-Palacin, and Danny Weyns. 2020. Understanding Uncertainty in Self-adaptive Systems. In *IEEE International Conference on Autonomic Computing and Self-Organizing Systems, ACSOS 2020, Washington, DC, USA, August 17-21, 2020*. IEEE, 242–251. doi:10.1109/ACSOS49614.2020.00047

[5] Winifred X. Chen, Nicholas M. Houghton, and Barrett S. Caldwell. 2022. Reducing the "Fog of Uncertainty" Surrounding Humanitarian Aid and Disaster Response Operations. *Advances in Transdisciplinary Engineering* 28 (2022), 144–153. doi:10.3233/ATDE220641 Transdisciplinarity and the Future of Engineering.

[6] Hugo Cruz-Suárez, Raúl Montes-de Oca, and R. Israel Ortega-Gutiérrez. 2023. Deterministic Discounted Markov Decision Processes with Fuzzy Rewards/Costs. *Fuzzy Information and Engineering* 15, 3 (2023), 274–290. doi:10.26599/FIE.2023.9270020

[7] Ali Goli and Behnam Malmir. 2020. A Covering Tour Approach for Disaster Relief Locating and Routing with Fuzzy Demand. *International Journal of Intelligent Transportation Systems Research* 18 (2020), 140–152. doi:10.1007/s13177-019-00185-2

[8] Vincenzo Grassi, Raffaela Mirandola, and Diego Perez-Palacin. 2025. Embracing Uncertainty: a Fuzzy Theoretical Model for Goal Fulfillment Assessment. In *2025 IEEE/ACM 13th International Conference on Formal Methods in Software Engineering (FormaliSE)*. 113–123. doi:10.1109/FormaliSE66629.2025.00018

[9] Mahbubeh Haghi, Seyed Mohammad Taghi Fatemi Ghomi, and Fariborz Jolai. 2017. Developing a robust multi-objective model for pre/post disaster times under uncertainty in demand and resource. *Journal of Cleaner Production* 154 (2017), 188–202. doi:10.1016/j.jclepro.2017.03.102

[10] Arnd Hartmanns, Sebastian Junges, Tim Quatmann, and Maximilian Weininger. 2023. A Practitioner's Guide to MDP Model Checking Algorithms. In *Tools and Algorithms for the Construction and Analysis of Systems*, Sriram Sankaranarayanan and Natasha Sharygina (Eds.). Springer Nature Switzerland, Cham, 469–488.

[11] Christian Hensel, Sebastian Junges, Joost-Pieter Katoen, et al. 2022. The probabilistic model checker Storm. *International Journal on Software Tools for Technology Transfer* 24, 4 (2022), 589–610. doi:10.1007/s10009-021-00633-z

[12] Sohaiba Iqbal, Muhammad Usama Sardar, Faiq Khalid Lodhi, and Osman Hasan. 2018. Statistical model checking of relief supply location and distribution in natural disaster management. *International Journal of Disaster Risk Reduction* 31 (2018), 1043–1053. doi:10.1016/j.ijdrr.2018.04.010

[13] Kenneth Johnson, Javier Cámara, Roopak Sinha, Samaneh Madanian, and Dave Parry. 2021. Towards Self-Adaptive Disaster Management Systems. In *Proceedings of the 18th ISCRAM Conference*. 49–61.

[14] Kenneth Johnson and Meena Kumari. 2025. Probabilistic Model Checking of Disaster Resource Distribution Strategies. In *2025 IEEE/ACM 13th International Conference on Formal Methods in Software Engineering (FormaliSE)*. 1–5. doi:10.1109/FormaliSE66629.2025.00012

[15] Kenneth Johnson, Samaneh Madanian, and Roopak Sinha. 2020. Graph-Theoretic Models of Resource Distribution for Cyber-Physical Systems of Disaster-Affected Regions. (8 2020), 521–528. doi:10.1109/SEAA51224.2020.00087

[16] Emine Nisa Kapukaya and Sule Itir Satoglu. 2025. A Multi-Objective Dynamic Resource Allocation Model for Search and Rescue and First Aid Tasks in Disaster Response by Employing Volunteers. *Logistics* 9, 1 (2025). doi:10.3390/logistics9010041

[17] George J. Klir and Bo Yuan. 1995. *Fuzzy sets and fuzzy logic: theory and applications*. Upper Saddle River, NJ: Prentice Hall.

[18] Masami Kurano, Masami Yasuda, Jun ichi Nakagami, and Yuji Yoshida. 1996. Markov-type fuzzy decision processes with a discounted reward on a closed interval. *European Journal of Operational Research* 92, 3 (1996), 649–662. doi:10.1016/0377-2217(95)00140-9

[19] M. Kurano, M. Yasuda, J. Nakagami, and Y. Yoshida. 2006. A fuzzy approach to Markov decision processes with uncertain transition probabilities. *Fuzzy Sets and Systems* 157, 19 (2006), 2674–2682. doi:10.1016/j.fss.2004.10.023 Fuzzy Sets and Probability/Statistics Theories.

[20] M. Kwiatkowska, G. Norman, and D. Parker. 2011. PRISM 4.0: Verification of Probabilistic Real-time Systems. In *Proc. 23rd International Conference on Computer Aided Verification (CAV'11) (LNCS, Vol. 6806)*, G. Gopalakrishnan and S. Qadeer (Eds.). Springer, 585–591.

[21] Samaneh Madanian, Kenneth Johnson, Mathew St Martin, Roopak Sinha, Javier Cámara, and David Parry. 2022. Adaptable socio-cyber physical systems for supporting disaster response. *Australasian Journal of Disaster and Trauma Studies* 26 (2022), 221–234.

[22] Sara Mahdavi-Hezavehi, Danny Weyns, Paris Avgeriou, Radu Calinescu, Raffaela Mirandola, and Diego Perez-Palacin. 2020. Uncertainty in Self-adaptive Systems: A Research Community Perspective. *ACM Trans. Auton. Adapt. Syst.* 15, 4 (2020), 10:1–10:36. doi:10.1145/3487921

[23] S. Nayeri, Z. Sazvar, and J. Heydari. 2022. A fuzzy robust planning model in the disaster management response phase under precedence constraints. *Operational Research: An International Journal* 22 (2022), 3571–3605. doi:10.1007/s12351-022-00694-1

[24] María José Pardo and David de la Fuente. 2010. Fuzzy Markovian decision processes: Application to queueing systems. *Computers & Mathematics with Applications* 60, 9 (2010), 2526–2535. doi:10.1016/j.camwa.2010.08.004

[25] Fatemeh Sabouhi, Ali Bozorgi-Amiri, and Parinaz Vaez. 2020. Stochastic optimization for transportation planning in disaster relief under disruption and uncertainty. *Kybernetes* 50, 9 (12 2020), 2632–2650. arXiv:https://www.emerald.com/k/article-pdf/50/9/2632/1751571/k-10-2020-0632.pdf doi:10.1108/K-10-2020-0632

[26] Ihab Hanna Sawalha. 2020. A contemporary perspective on the disaster management cycle. *Foresight* 22, 4 (05 2020), 469–482. arXiv:https://www.emerald.com/fs/article-pdf/22/4/469/614236/fs-11-2019-0097.pdf doi:10.1108/FS-11-2019-0097

[27] Lipika Shaw, Soumen Kumar Das, and Sankar Kumar Roy. 2022. Location-allocation problem for resource distribution under uncertainty in disaster relief operations. *Socio-Economic Planning Sciences* 82 (2022), 101232. doi:10.1016/j.seps.2022.101232

[28] H. Tanaka and K. Asai. 1984. Fuzzy linear programming problems with fuzzy numbers. *Fuzzy Sets and Systems* 13, 1 (1984), 1–10. doi:10.1016/0165-0114(84)90022-8

[29] Danny Weyns, Radu Calinescu, Raffaela Mirandola, Kenji Tei, and et al. 2023. Towards a Research Agenda for Understanding and Managing Uncertainty in Self-Adaptive Systems. *ACM SIGSOFT Softw. Eng. Notes* 48, 4 (2023), 20–36. doi:10.1145/3617946.3617951

[30] Yuji Yoshida. 2021. Markov Decision Processes with Fuzzy Risk-Sensitive Rewards: The Best Coherent Risk Measures Under Risk Averse Utilities. In *Computational Intelligence*, Juan Julián Merelo, Jonathan Garibaldi, Alejandro Linares-Barranco, Kevin Warwick, and Kurosh Madani (Eds.). Springer International Publishing, Cham, 135–161.