The Creation of Puffin, the Automatic Uncertainty Compiler

Nicholas Gray[a,b,*], Marco de Angelis[a], Scott Ferson[a]

[a]*Institute for Risk and Uncertainty, University of Liverpool, Liverpool, United Kingdom, L69 7ZX*
[b]*nickgray@liverpool.ac.uk*

Abstract

An uncertainty compiler is a tool that automatically translates original computer source code lacking explicit uncertainty analysis into code containing appropriate uncertainty representations and uncertainty propagation algorithms. We have developed an prototype uncertainty compiler along with an associated object-oriented uncertainty language in the form of a stand-alone Python library. It handles the specifications of input uncertainties and inserts calls to intrusive uncertainty quantification algorithms in the library. The uncertainty compiler can apply intrusive uncertainty propagation methods to codes or parts of codes and therefore more comprehensively and flexibly address both epistemic and aleatory uncertainties.

Keywords: Uncertainty Analysis; Uncertainty Compiler; Probability Bounds Analysis

1. Introduction

Modern science and engineering is all about numerical calculation. With the inexorable growth of computer power, more of these calculations are being undertaken with ever more complex computer simulations. These developments mean that new computation-intensive technologies are being explored, such as digital twins (see [1, Sec. 2.2.3.3] or [2]). Scientists and engineers need to make calculations even when there is uncertainty about the quantities involved, yet the tools they are commonly using do not allow this to be done intrusively. As a result many analysts work with computer codes that do not take full account of uncertainties.

Within the numerical calculations essential to engineering, there are two types of uncertainty: aleatory and epistemic. Aleatory uncertainty arises from the natural variability in changing environments and material properties, errors in manufacturing processes or inconsistencies in the realisations of systems. Aleatory uncertainty cannot be reduced by empirical effort. Epistemic uncertainty is caused by measurement imperfections or lack of perfect knowledge about a system. This could be due to not knowing the full specification of a system in the early phases of engineering design or ignorance about the expected manufacturing variations or deployment conditions. Imperfect scientific understanding of the underlying physics or biology involved causes uncertainty in predictions about the future performance of a system even after the design specifications have been decided. If uncertainties are small they can often be neglected or swept away by looking at the worst-case scenarios. However, in situations where the uncertainty is large, this approach is suboptimal or impossible, especially if it would impact a decision. Instead, a strategy of comprehensively accounting for the two kinds of uncertainty is needed that can propagate imprecise and variable numerical information through calculations.

Because analysts are typically unwilling to rewrite their codes, various simple strategies have been used to remedy the problem, such as elaborate sensitivity studies or wrapping the program in a Monte Carlo loop. These approaches treat the program like a black box because users consider it uneditable. However, whenever it is possible to look inside the source code, it is better characterised as a crystal box because the operations involved are clear but fixed and unchangeable in the mind of the current user.

---

[*]Corresponding author

2. The Problem with Monte Carlo

The most common approach to deal with uncertainty is to wrap code within a Monte Carlo shell. In this approach the calculations are repeated with random values for selected input variables. This is done for a large number of iterations, and the distribution of resulting outputs can be analysed. Such tools exists in many programming langages: DAKOTA for C++ [3], COSSAN [4] and UQLab [5] for MATLAB or UQpy for Python [6]. Olivier et al. give an excellent overview of many more software packages that are availible for non-intrusive uncertainty quantification [6]. Under such an approach random values are chosen and then the calculations are performed and the output stored, this is done for a number of iterations and total outputs can be analysed after the process has been completed.

In order to demonstrate the potential problems with such an approach we can consider a simple example. Suppose we have five variables $x_1, \ldots, x_5$ which are known to all have a value between 0 and 1 but no further information is known about the values. Suppose we need to perform the calculation

$$y = x_1 + x_2 + x_3 + x_4 + x_5 \tag{1}$$

with the knowledge that some bad thing will happen if $y \geq 4.5$. A number can be randomly generated for x_1, x_2, etc and these can be used in order to calculate the value of y for N iterations. After this is complete we can plot a histogram to show the distribution for y. Since we do not have any information about the distribution for $x_1, \ldots, x_5$ it seems sensible to assume that all values are equally likely and use a uniform distribution. Figure 1 shows these histograms for various N. From this we can see that as $N \to \infty$ the histogram resembles a normal distribution.

Whatever the number of replications used in the Monte Carlo simulation, we can estimate that the probability of the bad thing happening. With 10^6 replications, this estimate is $\Pr(y \geq 4.5) = 2.53 \times 10^{-4}$. However, it seems reasonable to consider whether we have confidence that the event is so rare. We had no information about the distributions of the five values except that they were between 0 and 1. Nor did we have knowledge about what dependencies there might be between the variables. From this information we cannot rule out the possibility that each x value is much more likely to be closer to 1 than 0, or that there is some dependence between the x values such that if x_1 is high then all the others are also likely to be high. Thus, the way that the uncertainty has been characterised may be significantly underestimating the risk [7].

There have been several engineering failures that were due in part to underestimating risks in ways similar to this example [7, 8]. Before the 1986 Challenger Disaster, NASA management had predicted the probability of failure with loss of vehicle and crew as 1 in 10^5 flights [9]. This turned out to be a gross underestimation of the true risk, which after the retirement of the fleet stood at 2 in 135. The Fukushima Daiichi nuclear disaster was due in part to underestimating the risk of a tsunami of the magnitude that caused the disaster and in failing to understand that collocating the backup generators created dependence that destroyed the planned engineered redundancy when the site was flooded during the event [10, p. 48]. The probabilities of satellites colliding in orbit can be underestimated through the use of probabilities [11], leading to false confidence that they are not going to hit each other.

Performing uncertainty analysis by simply wrapping a simulation code in a Monte Carlo loop may not give a full account of the uncertainties that are present within a simulation. The probabilities of extreme events are especially difficult to correctly estimate when either the distributions of input variables are not known or any inter-variable dependencies are not known. There are other limitations of this simplistic Monte Carlo approach, including false confidence [11], and problems arising from confounding epistemic and aleatory uncertainties [12].

3. Puffin

Strategies are needed that automatically translate original source code into code with appropriate uncertainty representations and propagation algorithms. Perez et al. introduced a MATLAB toolbox to perform automatic uncertainty propagation using unscented transform, however more general approaches are needed [13]. In this paper
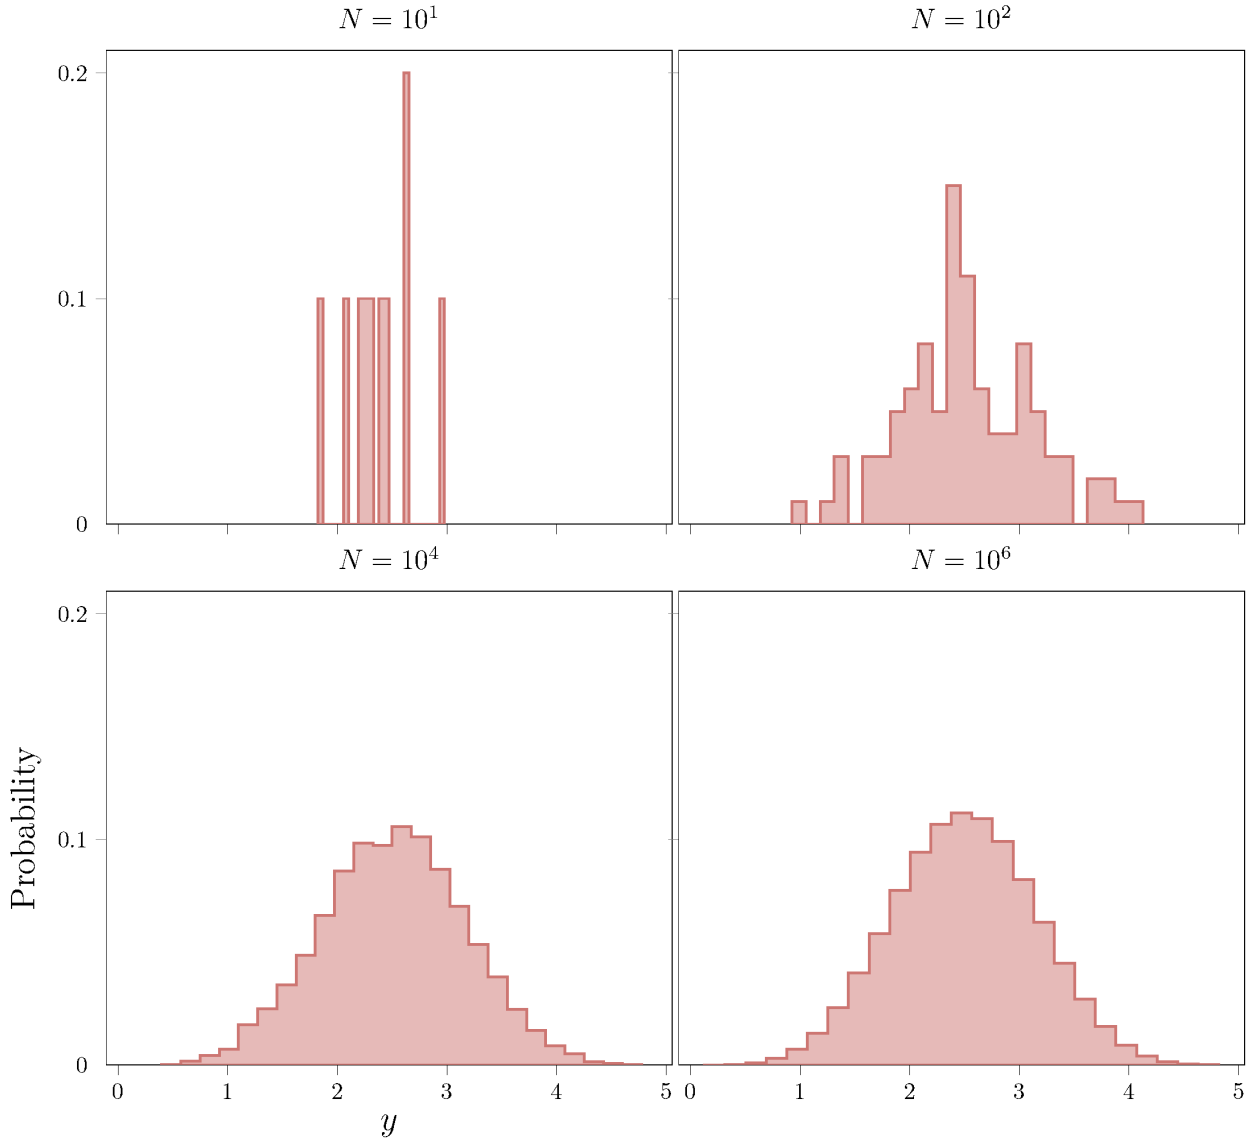
Fig. 1. Normalised histogram for the Monte Carlo simulation of Equation 1 for increasing number of iterations.

we describe an uncertainty compiler for this purpose, named Puffin, along with an associated language. It handles the specifications of input uncertainties and inserts calls to an object-oriented library of intrusive uncertainty quantification (UQ) algorithms. In theory, the approach could work with any computer language and any flavour of uncertainty propogation. There are several components that are needed for the creation of Puffin as shown in Figure 2.

Puffin needs a language of uncertainty, "Puffin Language", that allows users to specify what uncertainties should be associated with the variables within the source code. This language should be simple and independent of the source language. For every source language that Puffin is to work with, there has to be an intrusive UQ library that Puffin Language can be translated into. Puffin Language does not need to be a textual programming language, it could instead be a visual language as part of a graphical user interface. Every type of uncertainty that is expressible in Puffin Language must be supported by an object constructor in the uncertainty library written in the source language. Section 4 discusses this Puffin Language component.

The other component that Puffin needs is a transcompiler [14] that translates a user's source code into an UQ enriched code expressed in the same language. There are three subcomponents to this: a Reader that is able to read the source language, a Translator that is able to read the specified uncertainties in Puffin Language, and a Writer that
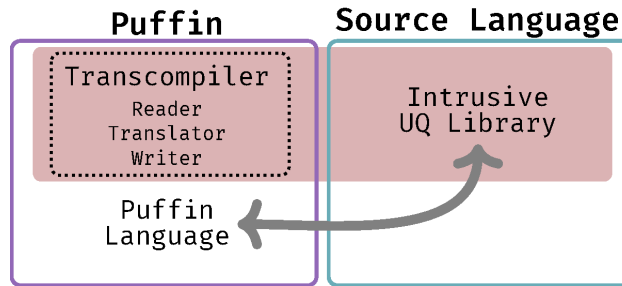
Fig. 2. The different components of Puffin and the UQ library it depends on. For each source language that Puffin is able to read, the parts highlighted in red need to be created. The intrusive uncertainty quantification library and Puffin Language need to mirror each other with a direct translation being available for all uncertainty specifications.

combines the results of both to output a new script with the specified uncertainties. ANTLR, a parser/lexer generator, can be used to generate the Reader and Writer [15]. ANTLR requires a grammar specification for the source language. Fortunately, ANTLR grammar files have been defined for many popular programming languages [16] and these can be used as a starting point. The Reader scans the input script and identifies the assignment operators within that may have uncertainties that need to be specified. The Translator reads Puffin Language specifications created by the user and translates the uncertainties into the source language. For each source language these conversions need to be specified. The Writer reproduces the script with the required uncertainties and necessarily changes required for the analysis to run without issue. We have designed these components for Python.

4. Puffin Language

Puffin depends on an uncertainty language to express what uncertainties are present within their scripts. This language enables users to specify the uncertainties about the variables involved in their code before compiling it into a new script with UQ enriched code. The language currently enables calculations with five types of uncertain objects that have relevance in engineering [12]:

- Interval: unknown value or values for which sure bounds are known [17],

- Probability distribution: random values varying according to specified law such as normal, uniform, binomial, etc., with known parameters [18],

- Probability box: random values for which the probability distribution cannot be specified exactly but can be bounded [19],

- Confidence structure: inferential uncertainty about a constant parameter compatible with both Bayesian and frequentest paradigms [20], and

- Natural language expressions: uncertain values indicated by linguistic hedge words such as 'about 7.2' or 'at most 12' [21, 22, 23].

Libraries that add some/all of these objects are available in C++ [24], Python [25], MATLAB [26], R [27 or 28], Julia [29]. There are many other uncertain objects that could be included within such a language of uncertainty such as second-order distributions or meta-distributions [30, 31], fuzzy numbers [32], possibility distributions [33], consonant structures [34], info-gap models [35] and others.

4.1. Intervals

An interval is an uncertain number representing values from an unknown distribution over a specified range, or perhaps a single value that is imprecisely known even though it may in fact be fixed and unchanging. Intervals thus

embody epistemic uncertainty. Intervals can be specified by a pair of scalars corresponding to the lower and upper bounds of the interval, such as $[0, 1]$ or $[4, 5]$. They can also be expressed as a value plus or minus some error, such as $[5 \pm 2]$ which is equivalent to $[3, 7]$.

Interval arithmetic computes with ranges of possible values, as if many separate calculations were made under different scenarios. However, the actual computations made by the software are done all at once, so they are very efficient. Basic binary operations $(+, -, \times, \div)$ can be performed using interval arithmetic:

$$[a, b] + [c, d] = [a + c, b + d], \tag{2}$$

$$[a, b] - [c, d] = [a - d, b - c], \tag{3}$$

$$[a, b] \times [c, d] = [\min(ac, ad, bc, bd), \max(ac, ad, bc, bd)] \tag{4}$$

and

$$[a, b] \div [c, d] = \left[\min\left(\frac{a}{c}, \frac{a}{d}, \frac{b}{c}, \frac{b}{d}\right), \max\left(\frac{a}{c}, \frac{a}{d}, \frac{b}{c}, \frac{b}{d}\right)\right] \tag{5}$$

assuming that $0 \notin [c, d]$.

Intervals can be propagated through all common mathematical functions such as exp, sin, log, etc. This is relatively straightforward if the function is monotonic as this implies that the endpoints of the input interval correspond to the endpoints of the output interval. For example, when calculating the exponential of an interval,

$$\exp([0, 1]) = [\exp(0), \exp(1)] \approx [1, 2.718] \tag{6}$$

For a non-monotonic function such as sine it is not necessarily the case that the endpoints of the interval correspond the the endpoints of the output function. For example, it is not the case that

$$\sin([0, \pi]) = [\sin(0), \sin(\pi)] = [0, 0] \tag{7}$$

because there are many x values such that $\sin(x) > 0$ for $x \in [0, \pi]$. The true width of the interval can be calculated using the maxima of the function within the domain, in this case $\frac{\pi}{2}$. Hence,

$$\sin([0, \pi]) = [0, 1]. \tag{8}$$

4.2. Probability Distributions and Probability Boxes

A probability distribution is a mathematical function that gives the probabilities of occurrence for different possible values of a random variable.

Probability bounds analysis integrates interval analysis and probability distributions using probability boxes (p-boxes) [19]. They can be considered as interval bounds on a probability distribution, therefore one can think of a probability distribution as a special case of a p-box. Figure 3 shows an example of a probability distribution. P-boxes characterise both epistemic and aleatory uncertainty. A p-box can be expressed mathematically as

$$\mathcal{F}(x) = [\underline{F}(x), \overline{F}(x)], \; \underline{F}(i) \geq \overline{F}(i) \; \forall x \in \mathbb{R} \tag{9}$$

where $\underline{F}(x)$ is the function that defines the left bound of the p-box (the blue in Figure 3) and $\overline{F}(x)$ defines that right bound of the p-box (the orange line in Figure 3)

As with intervals, standard arithmetic operations can be performed on p-boxes (and therefore probability distributions). For two p-boxes $\mathcal{A}(x) = [\underline{A}(x), \overline{A}(x)]$ and $\mathcal{B}(x) = [\underline{B}(x), \overline{B}(x)]$, then

$$\mathcal{C}(x) = \mathcal{A}(x) \circ \mathcal{B}(x) = [\underline{C}(x), \overline{C}(x)] \tag{10}$$
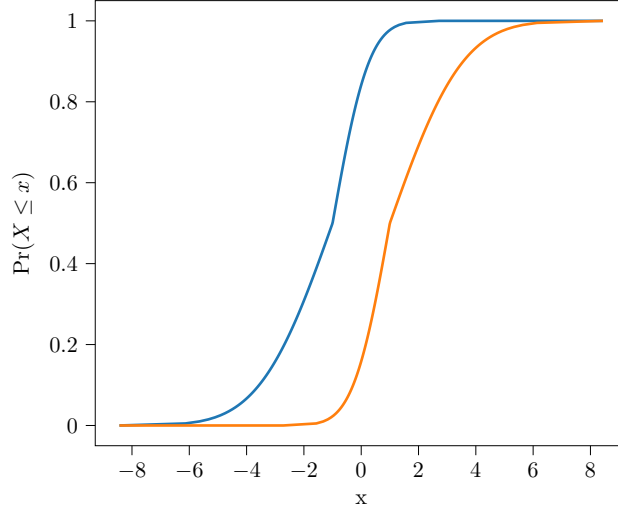
5

Fig. 3. Probability box for a normal distribution with $\mu = [-1, 1]$ and $\sigma = [1, 2]$.

where

$$\underline{C(z)} = \inf_{z=x \circ y} \left[\min \left(\underline{A(x)} \circ \underline{B(y)}, 1 \right) \right] \tag{11a}$$

$$\overline{C(z)} = \sup_{z=x \circ y} \left[\max \left(\overline{A(x)} \circ \underline{B(y)} - 1, 0 \right) \right] \tag{11b}$$

if $\circ \in [+, \times]$, or

$$\underline{C(z)} = 1 + \inf_{z=x \circ y} \left[\min \left(\underline{A(x)} \circ \overline{B(y)}, 0 \right) \right] \tag{12a}$$

$$\overline{C(z)} = \sup_{z=x \circ y} \left[\max \left(\overline{A(x)} \circ \underline{B(y)}, 0 \right) \right] \tag{12b}$$

if $\circ \in [-, \div]$. Naturally, division is only valid if $0 \notin \mathcal{B}$ [36, p. 89].

Within a programming language a p-box can be expressed by using the name of the probability distribution, or some shorthand for the name, as the function and the arguments as intervals. For example, the p-box shown in Figure 3 could be generated using

```
Normal([-1,1],[1,2]).
```

P-boxes can be defined in situations where the shape of the distribution is unknown but some empirical evidence about the data is known, such as the minimum, maximum, mean, standard deviation, etc. In this situation bounds can be created such that they are consistent with all the available information [37].

4.3. Confidence Boxes

Confidence boxes (c-boxes) are imprecise generalisations of traditional confidence distributions, which, like Student's t-distribution, encode frequentist confidence intervals for parameters of interest at every confidence level [20, 38]. They are analogous to Bayesian posterior distributions in that they characterise the inferential uncertainty about distribution parameters estimated from sparse or imprecise sample data, but they have a purely frequentist interpretation that makes them useful in engineering because they offer a guarantee of statistical performance through repeated use. Unlike confidence intervals which cannot usually be used in mathematical calculations, c-boxes can be propagated through mathematical expressions using the ordinary machinery of probability bounds analysis, and this allows analysts to compute with confidence, both figuratively and literally, because the results also have the same confidence interpretation [39]. For instance, they can be used to compute probability boxes for both prediction and tolerance distributions.

Confidence boxes can be computed in a variety of ways directly from random sample data. There are c-boxes both for parametric problems (where the family of the underlying distribution from which the data was randomly generated is known to be normal, binomial, Poisson, etc.), and for non-parametric problems in which the shape of the underlying distribution is unknown. C-boxes account for the uncertainty about a parameter that comes from the inference about observations, including the effect of small sample size, but also the effects of imprecision in the data and demographic uncertainty which arises from trying to characterise a continuous parameter from discrete data observations.

For example, it is possible to specify a c-box in the binomial case of having K successes out of N trials, based upon Clopper-Pearson confidence intervals [34, 40, 41]. This K-out-of-N c-box is specified as

$$\mathcal{KN}(k, n) = [\text{beta}(k, n - k + 1), \text{beta}(k + 1, n - k)] .$$ (13)

4.4. Natural Language Uncertainty

In order to make uncertainty analysis as simple as possible, users should be able to input their uncertainties using natural language expressions such as *about* or *almost*. Humans are more likely to express their uncertainties in terms of hedged expressions around a round number, rather than as a percentage or probability. Table 1 lists some hedge words and their possible interpretations. Hedge words can be interpreted as intervals, p-boxes [21], or consonant c-boxes [42].

Hedged Numerical Expression	Possible Interpretation
ABOUT(x)	$[x \pm 2 \times 10^{-d}]$
AROUND(x)	$[x \pm 10 \times 10^{-d}]$
COUNT(x)	$[x \pm \sqrt{x}]$
ALMOST(x)	$[x - 0.5 \times 10^{-d}, x]$
OVER(x)	$[x, x + 0.5 \times 10^{-d}]$
ABOVE(x)	$[x, x + 2 \times 10^{-d}]$
BELOW(x)	$[x - 2 \times 10^{-d}, x]$
AT MOST(x)	$[0, x]$
AT LEAST(x)	$[x, \infty]$
ORDER(x)	$[x/2, 5x]$
BETWEEN x AND y	$[x, y]$
K OUT OF N	$[\text{beta}(k + 1, n - k), \text{beta}(k, n - k + 1)]$

Table 1. Hedge expressions and their mathematical equivalent. Note: d is the number of significant figures of x.

4.5. Logical Operations with Uncertain Objects

When making decisions it is often the case that two values need to be compared with each other. Asking whether an observed value is greater than, equal to, or less than some threshold value is fundamental. For example, if a decision relies on some observed value X being less than 1, when we know the value of X accurately then it is easy to make such a comparison. However, if there is some uncertainty about the value of X then this comparison may not be so easy.

For intervals, $X = [a, b]$ and $Y = [c, d]$, then

$$X < Y = \begin{cases} 1 & b < c \\ 0 & a \geq c \\ [0, 1] & \text{otherwise} \end{cases}$$ (14)

and

$$X > Y = \begin{cases} 0 & b \leq c \\ 1 & a > d \\ [0, 1] & \text{otherwise} \end{cases}$$ (15)

with 0 and 1 denoting true and false respectively, and [0,1] being the Boolean equivalent of "I don't know". We can call [0,1] the *dunno* interval. This implies that we cannot say whether an uncertain value characterised by an interval is larger or smaller than another unless the interval is entirely greater or less than the other interval. For the equality comparison,

$$X == Y = \begin{cases} [0,1] & a \in Y \text{ or } b \in Y \\ 0 & \text{otherwise,} \end{cases} \tag{16}$$

when asking for equivalence between intervals it is never possible to say that one value is equal to another. We can introduce a new Boolean operator (===) to test for whether two uncertain numbers are equivalent in form,

$$X === Y = \begin{cases} 1 & a = c \text{ and } b = d \\ 0 & \text{otherwise.} \end{cases} \tag{17}$$

The dunno interval can be converted into a true Boolean using operators such as ALWAYS or SOMETIMES

$$\text{ALWAYS}\,([0,1]) = 0 \tag{18a}$$

$$\text{SOMETIMES}\,([0,1]) = 1 \tag{18b}$$

so that we can get

$$\text{ALWAYS}\,(X < Y) = \begin{cases} 1 & b < c \\ 0 & \text{otherwise} \end{cases} \tag{19}$$

$$\text{SOMETIMES}\,(X < Y) = \begin{cases} 1 & a < d \\ 0 & \text{otherwise.} \end{cases} \tag{20}$$

There are methods that are able to deal with more nuanced ways of using logical operations with intervals, see [43] as an example. There are also different logic systems such as fuzzy logic that could be used in order to make logical operations with uncertain numbers [44].

4.6. Repeated Variables and Dependency

When performing intrusive uncertainty analysis it would be ideal to always obtain best possible results that are guaranteed to bound the true value without overestimating the uncertainty. The uncertainty can be inflated or artifactually high if careful consideration of the dependence between, and repetition of, uncertain numbers is not undertaken. This problem appears to be ubiquitous to many, if not all, uncertainty calculi [12].

For example, if $a = [2,3]$ and $b = [4,5]$, then $a \times b = [8,15]$. However, if it were the case that a and b were oppositely dependent on each other, such that a low value of a is always matched with a high value of b, then $a \times b$ is the much narrower interval [10,12].

Repetition of variables can also artifactually inflate the amount of uncertainty present within the output. For example if $a = [1,2]$, $b = [-1,1]$ and $c = [3,4]$ then

$$\begin{aligned} ab + ac &= [1,2] \times [-1,1] + [1,2] \times [3,4] \\ &= [-2,2] + [3,8] \\ &= [1,10] \end{aligned} \tag{21}$$

but

$$\begin{aligned} a(b+c) &= [1,2]\,([-1,1] + [3,4]) \\ &= [1,2] \times [2,5] \\ &= [2,10] \end{aligned} \tag{22}$$
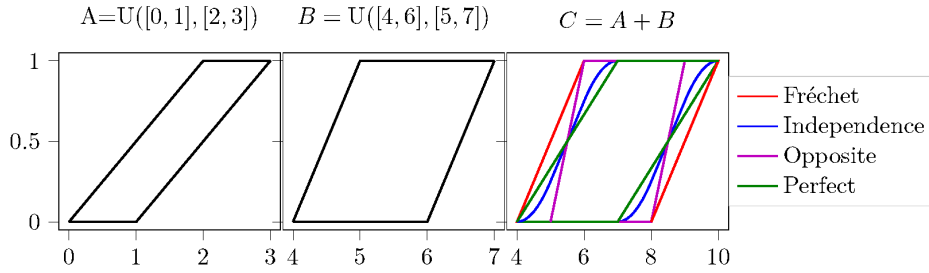
Fig. 4. Addition of two p-boxes with different dependencies.

Although algebraically these two expressions should be equal, the uncertainty of $ab + ac$ is greater than the uncertainty about $a(b + c)$. This is due to the fact that the uncertain variable a is repeated within the former but appears only once in the latter. In essence the uncertainty about a has been considered twice when performing the first calculation. The amount of this artifactual uncertainty can be reduced by transforming the original equation into a single-use expression where uncertain variables are only used once. If this is not possible, there are other techniques that can be used to reduce this artifactual uncertainty (e.g. [45, 46, 47, 48])

For distributions and p-boxes, significant artifactual uncertainty reduction can be made if the dependence between the variables is known [36]. Figure 4 shows the result of adding two separate p-boxes, $A = \mathrm{U}([0, 1], [2, 3])$, $B = \mathrm{U}([4, 6], [5, 7])$, together with different dependencies between A and B. The Fréchet bounds are used when the dependence between A and B is unknown, thus it is the most general case and is guaranteed to bound the correct answer. As such in Figure 4 the Fréchet bounds cover all the other dependencies, as it is the operation that is defined in equations 10, 11 and 12. Perfect, or comonotonic, is where there is perfect positive dependence between the two variables, with the highest possible correlation coefficient. Opposite, or countermonotonic, is perfect negative dependence between the two variable with the lowest possible correlation coefficient. Independence is where there is no dependence between the two variables. It should not be assumed that variables are independent unless this is known because wrongly assuming independence can lead to incorrectly reducing the amount of uncertainty and understating tail risks.

In general, dependence between uncertain quantities can be expressed through the use of correlation coefficients or copulas or bounds on copulas more generally [49, 50, 51, 52]. This can include named copulas such as independence, opposite and perfect, as shown in Figure 4, or other copula families parameterised by a numerical correlation coefficient. Independence implies the correlation is zero, although zero correlation does not imply independence. Likewise, a correlation of one implies perfect dependence, but, depending on the copula family, perfect dependence may not imply correlation one. The symbol $\equiv$ can be used to indicate that the variables are equal in value, i.e., equal in distribution and perfectly positively correlated.

These dependencies can be stored within a matrix, such as that shown in Table 2. This matrix can be checked for feasibility by checking that it is positive semi-definite and that there are no conflicting dependencies within the table. For example the matrix shown in Table 3 is not logically consistent for continuous variables. This is because a high value of x implies a high value of z, since they are positively dependent on each other. Meanwhile, a high z implies y must be low, due to their opposite dependence. This means there must also be dependence between x and y, not the independence as has been specified within the table.

4.7. Other issues

Aside from dependencies and sensitivity to repeated variables, there are other issues that distinguish simple deterministic calculations from uncertainty quantifications. For example, in uncertainty quantification analysts may need to consider *ensembles* and *backcalculations*.

Probability distributions describe properties or behaviours across a population of entities. Statisticians call such a population the "reference class" or "ensemble". Uncertainty quantification implicitly represents many calculations over

	a	b	c	d	e
a		i	≡	f	o
b	i		i	f	0
c	≡	i		f	o
d	f	f	f		f
e	o	0	o	f	

Table 2. Matrix showing dependencies between several variables. (f - Fréchet, i - Independence, o - Opposite, ≡ - Equal in value)

	x	y	z
x		i	p
y	i		o
z	p	o	

Table 3. Dependency matrix that does not make logical sense. (i - Independence, p - Perfect, o - Opposite)

interacting ensembles, and it can be extremely important to keep in mind what the values in a probability distribution represent. For instance, if the post-operative risks of prostatectomy is fifty percent erectile dysfunction, it would make a huge difference to a patient whether this means that 50% of his future attempts at sex fail or that 50% of patients are permanently impotent. Does a system reach 10% of criticality or does it reach criticality 10% of the time? Uncertainty quantifications that do not explicitly define what the distributions in an analysis represent in terms of their respective ensembles may be meaningless. Puffin allows users to annotate their codes to specify and document the ensemble described by any distribution or other uncertain quantity, although, in general, it is the responsibility of the analyst to ensure that the calculations used make sense.

Another wrinkle that makes uncertainty quantification different from its analogous deterministic calculations is the importance of backcalculation. Backcalculation is a mathematical operation for finding solutions to equations involving variability or uncertainty that guarantee some desired performance. Such problems are ubiquitous in engineering design. Backcalculation solves questions such as

1. What dimensional constraints on a component are necessary to ensure that it fits in its place in a machine given spatial tolerances?
2. How much propellant is needed to guarantee sufficient fuel given the mission contingencies and unforeseen variabilities?
3. How much shielding is needed on a spacecraft to ensure that the total ionizing radiation experienced inside the craft does not exceed some tolerable threshold, given that radiation in space varies over time in an imperfectly known way?

The Puffin UQ library has algorithms to solve backcalculations that involve intervals, distributions and p-boxes (when solutions exist), but it is the responsibility of the analyst to ensure that they are deployed appropriately to yield calculations that make sense in the engineering context.

5. Compiler

Puffin consists of its intrusive UQ library, a code inspector/editor, and an uncertainty compiler. Puffin's uncertainty compiler does five things:

1. Parses the input source code into expression tree(s),
2. Identifies the variables in any assignment operations,
3. Replaces or modifies some or all of these assignments in the expression trees according to options and specifications provided by the user,
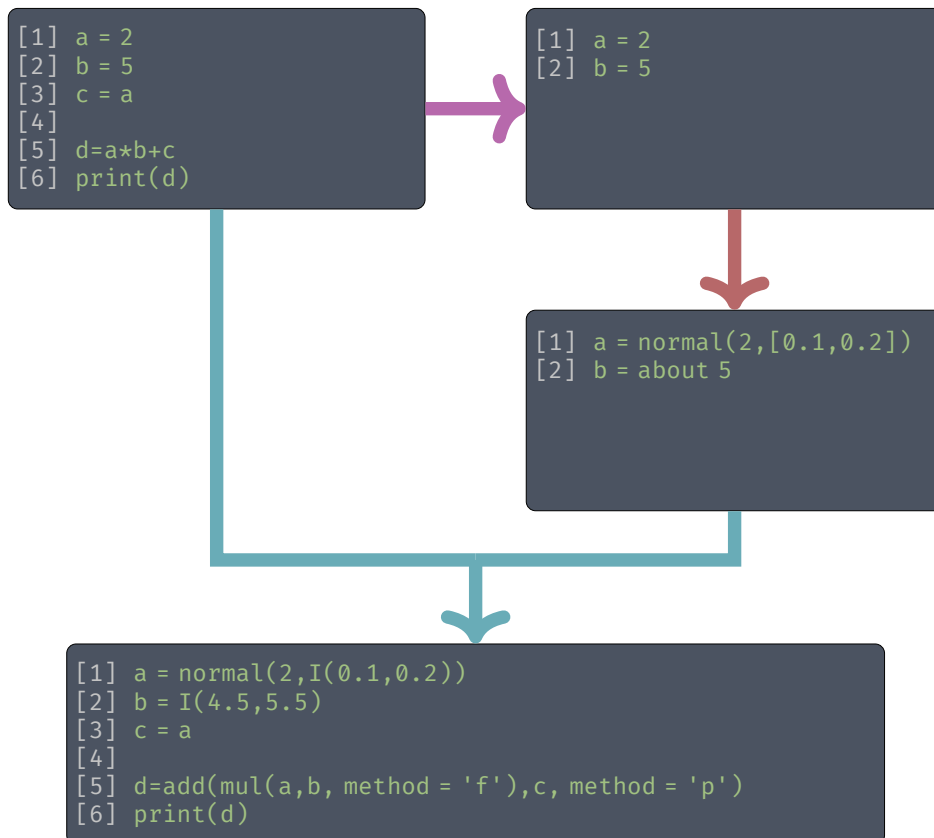
```
[1] a = 2                    [1] a = 2
[2] b = 5                    [2] b = 5
[3] c = a
[4]
[5] d=a*b+c
[6] print(d)


                             [1] a = normal(2,[0.1,0.2])
                             [2] b = about 5




[1] a = normal(2,I(0.1,0.2))
[2] b = I(4.5,5.5)
[3] c = a
[4]
[5] d=add(mul(a,b, method = 'f'),c, method = 'p')
[6] print(d)
```

Fig. 5. The result of using the compiler whilst defining the uncertainty in Puffin on a simple pseudocode script.

4. Translates the expression trees, with amended assignments, into the target language equipped with its intrusive UQ library, and

5. Analyses the output code to detect repeated variables and other functional dependencies that affect calculations and suggests improvements for computing uncertainties.

In order to explain what is happening in these steps it is useful to consider a simple pseudocode script, shown in the top left corner of Figure 5. For step 1, the simple script has then been broken into a parse tree which can be seen in Figure 6. From this tree Puffin detects the assignment operators which define a variable. These include lines 1 and 2, the leaves highlighted magenta on the parse tree, but not those that assign a value based upon a mathematical expression (line 5), a function or directly from another variable (line 3). In theory, such variables could also be edited by the user, but Puffin assumes that only explicit assignments will have uncertainty. Once the assignments have been found they can be displayed in the Puffin language, shown in the top right panel of Figure 5 where the assignments can then be edited with the appropriate user-specified uncertainties. These uncertainties need to then be translated to the source language, along with the rest of the parse tree.

This translation may include altering any functions that depend the amended variables. In this case, the infix operators in the definition of d in line 5 (+, *) have been identified within the parse tree and replaced with an explicit call to the UQ library functions (add, mul) which also have as an argument the dependence operation that is to be used. The lower panel of Figure 5, the value 'f' of this argument corresponds to making no assumption about the intervariable dependence between a and b, and perfect dependence (comonotonity) between their product and the variable c.

Puffin should only highlight numeric objects, not characters, strings or other non-numeric classes. In strongly typed programming languages like C, FORTRAN, and Pascal, the problem of distinguishing numeric from other types of objects is easy. In Python, R or Julia, the type of any object is not detectable until runtime and can even change during execution. Puffin will also need to be able to recognise objects that are collections of numeric values such as lists or
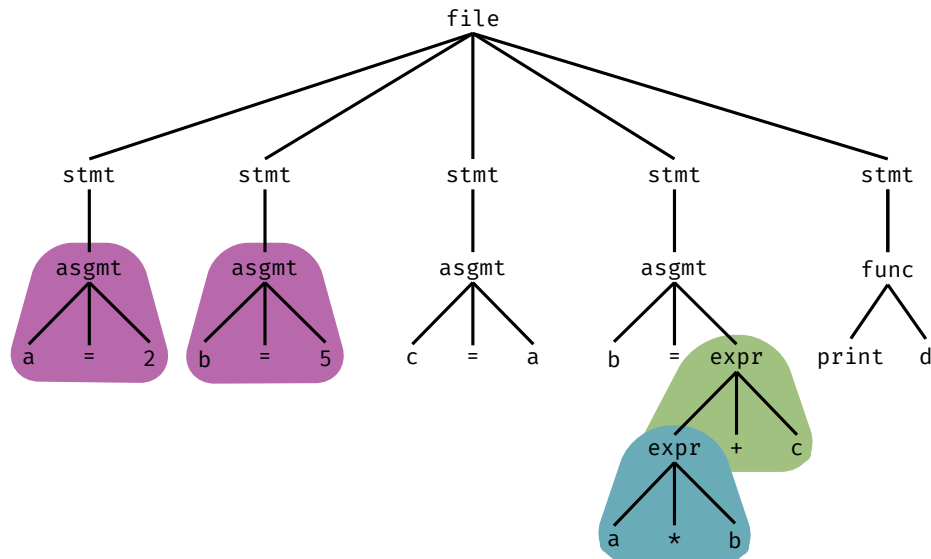
11

Fig. 6. Parse tree for the simple pseudocode script. Abbreviations: stmt - statement, asgmt - assignment, func - function, expr - expression.
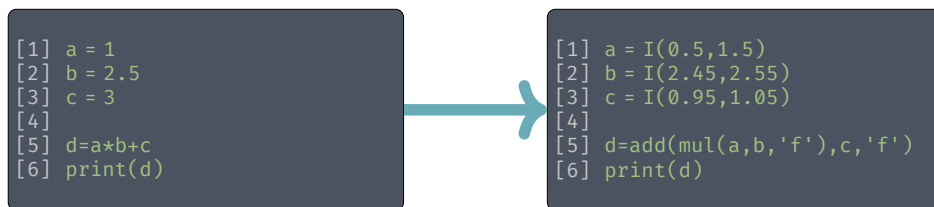

Fig. 7. Result of using Puffin automatically on a simple pseudocode script.

dictionaries. It will also need to be able to look at what is inside the lists and highlight those which have numeric objects within the list and allow of the individual objects to have uncertainties added. Alternatively it could be the case that the whole of the list has the same uncertainty, something which should be possible.

Puffin can be run automatically without any user input at all. Under default settings, automatic uncertainty compilation replaces floating-point constants with intervals interpreted from the significant figures used in the source code assignments and uses that information as a proxy for the uncertainty (for an example see Figure 7). In this mode all the steps of the steps happen concurrently without requiring any further input from an end user. When using this mode the compiler will need to tread carefully around mathematical constants such as π or e for which there is no uncertainty. Ideally it would allow users to minimally specify what values are precise constants.

5.1. Control Flow and Functions

For loops and functions here are potential stumbling blocks for Puffin. Figure 8 shows a simple pseudocode script with a function and a for loop. The first for loop each `i` is simply just a control variable with start and end variable. The individual value of `i` is irrelevant and as such would have no uncertainty about it. The second for loop is a 'for each' loop implying that the code needs to do something for each value within some iterable object, under this scenario it may be the case that there is uncertainty about the object within the list. In this case the code is setting the value of `initial_velocity` as each value within the the list, for each iteration of the for loop. In this case it may be the case that there is uncertainty within the object, in which case Puffin should recognise this and all users to change the code such that the objects within the list can have uncertainty added to them.

Puffin will also need to have a way of dealing with local variables within functions. For example, the function in Figure 8 has two local variables, `s` for the distance that the object travels and `g` for the acceleration due to gravity. It is conceivable that both of these variables have some uncertainty associated with them and as such Puffin should be able

```
[ 1] def calculateVelocity(u):
[ 2]     s = 10
[ 2]     g = 9.81
[ 3]     v = sqrt(u^2 + 2*g*s)
[ 4]     return v
[ 5]
[ 6] x = 1
[ 7] for i in (1:10):
[ 8]     x = x + 1
[ 9]
[10] for initial_velocity in (0.3,0.4,1.1,2.3,3.7):
[11]     final_velocity = calculateVelocity(initial_velocity)
[12]     print(final_velocity)
[13]
[14] a = 1
[15] b = 0.5
[16] if a < b:
[17]     c = 1
[18] else:
[19]     c = 2
[20]
```

Fig. 8. Pseudocode script with functions and for loops.

to detect the variables and offer the ability to edit them so that uncertainty is handled. This could be done using a dot notation, meaning that the local `g` can be accessed using `calculateVelocity.g`.

If statements and other logical control structures may also pose issues for Puffin. In line 16 of Figure 8 the logical operation within the statement would need changing to ensure that the statement runs as expected, see Section 4.5. Ideally, the analyst would decide what should happen if the statement `a < b` returned a dunno, [0,1], result by using the adverb operators discussed in equations 19 and 20. This may require additional editing to deal with situations where an uncertain result should be handled differently to a certain true or false.

All the code and variables are not necessarily contained within one script. For example, classes and functions are often placed in other files in order to improve readability or to avoid repetitions. Ideally Puffin would be able to parse several scripts at the same time and remember the context for all the individual objects. It is also often the case that scripts read data from other files when running. Under this scenario it may be difficult to use Puffin to express the uncertainty directly within the script, although import functions could be modified to add in the uncertainties. For instance, anytime a floating-point number is read from the file, its significant digits could be interpreted to specify an interval around the value. So, for example, the value '3.56' would be understood as the interval [3.555, 3.565]. Another approach might be to get Puffin to parse the data file and add the uncertainties in to the file directly. This would require changing the import function to be able to handle uncertain datafiles.

Many computers languages that are not purely functional support functions that specify their parameters with "call by reference" which means that the memory location of a value is passed to the function rather than a copy of the true value. This convention can allow the function to change the values of those parameters in the calling routine not just locally within the function. Python does this by default with objects more complicated than integers, floating-point number, and strings such as lists, dataframes and numpy arrays. Puffin will need to be careful in handling functions that use the call by reference method of passing argument.

The presence of uncertainty implies multiple function definitions might be useful. For instance `sqrt` applied to ranges that might include negative numbers could have three possible behaviours. Abnormal termination, for example Python's `math.sqrt` returns a domain error if passed a negative number; yielding imaginary results, such as Python's
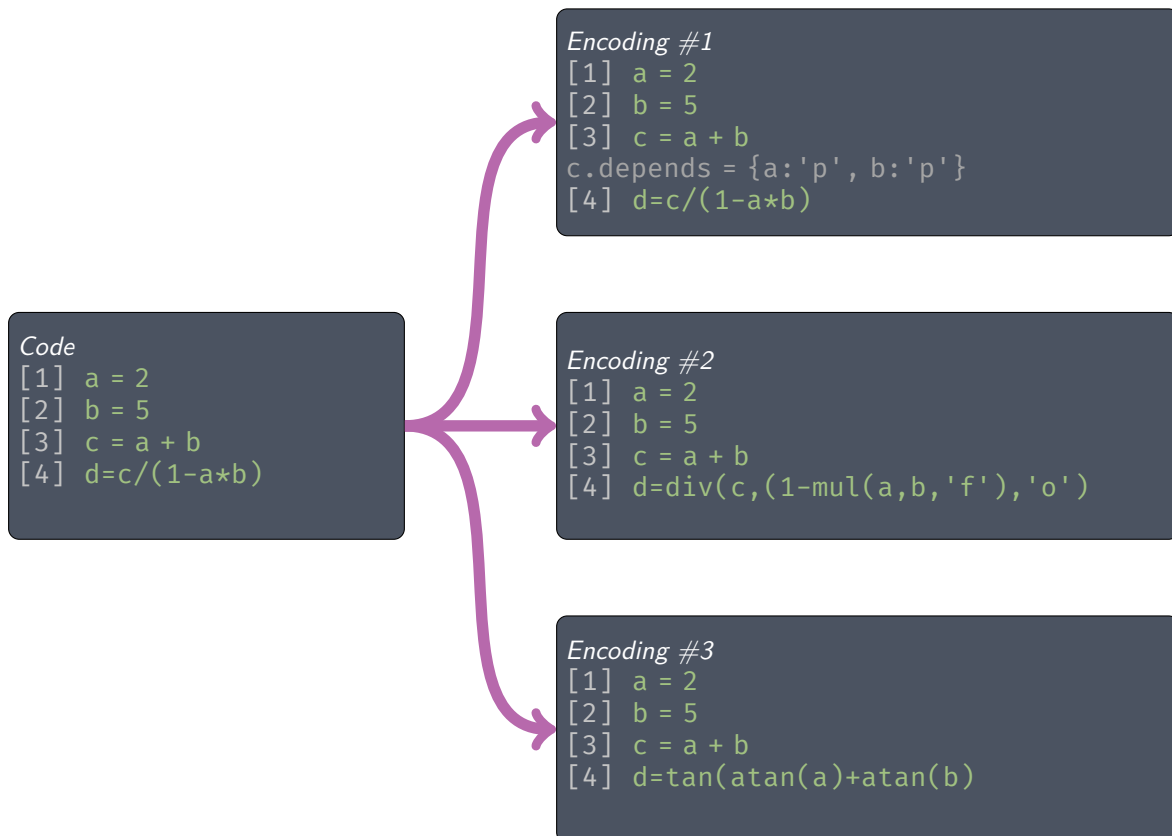
Fig. 9. Three possible encodings for the dependencies in the code shown on the left.

`cmath.sqrt`; or ignoring the negative values, returning [0,1] for $\sqrt{[-1, 1]}$[1]. An example of this can be seen when using `numpy.sqrt([-1,1])` in Python which returns [NaN, 1]. Similar facilities already exist to handle NaN's or missing values.

5.2. Coping with Dependency and Repeated Variable Problems

When it comes to dealing with the issues of dependency and repeated variables there are a couple of approaches that could be used in order to help reduced the problems discussed in Section 4.6. The simplest approach from a Puffin perspective would be for the libraries within each language to be able to handle the dependencies directly. This could be done if each object kept track of what other objects it depends on and in what way. For the example that has been used in Figure 9–Encoding #1, the variable `c` would remember that it is dependent on the variables `a` and `b` and therefore on line 4 it would know what the correct arithmetic would be to ensure as little artifactual uncertainty as possible. Puffin would insert this dependence directly in the translation as shown with the grey text in the Encoding #1, then at run time the / and ∗operators would automatically invoke the correct algorithms that respect the dependencies between the variables detected. This approach would have demands on memory. It requires initialised variables to have dependencies specified or a default dependence if they are unspecified.

The other way of treating the dependence would be for Puffin to parse over the script in order to detect the dependencies directly at compile time. These dependencies can then be stored within a matrix as discussed in Section 4.6. This matrix would need to be accessible for an analyst to add in assumptions about dependencies not observable from the code. For example, if variables `a` and `b` are independent this cannot be directly inferred from the code and therefore Fréchet would be assumed unless the analyst stated otherwise by editing the matrix. In this

---

[1] In this instance the square brackets correspond to Python's list object not an interval

scenario Puffin would have to replace any infix operators with function calls with the appropriate dependence. In Figure 9–Encoding #2, the multiplication a∗b has been replaced by the multiply function specifying that Fréchet should be used for the dependencies. Similarly the division operator has been replaced by a function with the method defined as opposite.

Finally where it is possible Puffin should be able to rearrange the equations such that any repeated variables are removed. Such a method would require Puffin to have a directory of multi-use to single-use rearrangements as well as a way of matching the written code to the mathematical expression. Another, smarter, approach would be to have a symbolic algebra system that is able to rearrange to a single-use expression on the fly. The simplest version of this is for it to happen just across one line, for instance replacing

```
c = a*b+a*c
```

with

```
c = a*(b+c).
```

A more complex approach is to consider repetitions globally to try to reduce the repetitions by detecting repetitions that happen over multiple lines. For instance in the example code in Figure 9, there is a hidden repetition across lines 3 and 4 because $c = a + b$. So

$$d = \frac{c}{1 - a * b} = \frac{a + b}{1 - a * b}. \tag{23}$$

This expression can be rearranged into the single-use expression

$$d = \tan\left(\arctan(a) + \arctan(b)\right) \tag{24}$$

which is the change made in Encoding #3. Such a transformation would be in the directory mentioned above.

Care would need to be taken to ensure that the right variable gets the rearrangement. Take the following kinematics equation to find the position s of a particle at time t

$$s = ut + \frac{1}{2}at^2 \tag{25}$$

where u is the initial velocity and a is the acceleration of the particle. This equation has a single repetition for both u and a but t is repeated. If there is uncertainty associated with t then this equation can be rearranged into a single-use expression

$$s = \left(\sqrt{\frac{a}{2}}t + \frac{u}{\sqrt{2a}}\right)^2 - \frac{u^2}{2a}. \tag{26}$$

This equation contains repetitions of a and u and as such may only be preferred if there is no uncertainty associated with either a or u. If there is uncertainty associated with either then it may be best not to perform the rearrangement or to intersect possible rearrangements to obtain the best possible expression.

There are additional issues that Puffin could face when rearranging equations. For example if the uncertainty about a includes negative numbers then $\sqrt{2a}$ is likely to be problematic. Alternatively, there could be problems if a straddles 0 because this would result in a division by zero. A strategy for dealing with this may be to perform the calculations in using both Equation 25 and 26 and intersect them.

5.3. Hermeneutic Problems

There are several problems that could occur when it comes to translating a script because it is difficult to understand the intent of the programmer from the code.

An example of this can be found in line 3 of the psuedocode example in Figure 5. There is potential for confusion when it comes to the assignment c = a as there are a couple of different interpretations as to what such a command implies when it comes to the uncertainty. The first is that we are implying that c and a are the same object but have
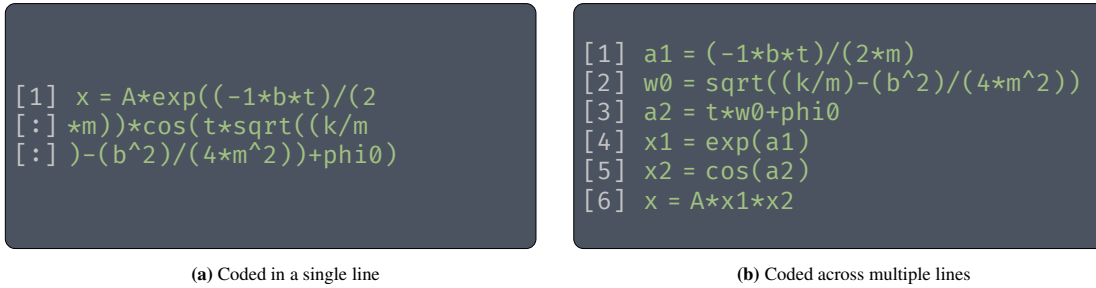
```
[1] x = A*exp((-1*b*t)/(2
[:] *m))*cos(t*sqrt((k/m
[:] )-(b^2)/(4*m^2))+phi0)
```

(a) Coded in a single line

```
[1] a1 = (-1*b*t)/(2*m)
[2] w0 = sqrt((k/m)-(b^2)/(4*m^2))
[3] a2 = t*w0+phi0
[4] x1 = exp(a1)
[5] x2 = cos(a2)
[6] x = A*x1*x2
```

(b) Coded across multiple lines

Fig. 10. Two different way in which Equation 28 might be coded.

been given different names for some reason, under this scenario they should be considered equivalent to each other and therefore the calculation a + c could be rearranged to the single use expression 2*a. A second interpretation is to consider that the line could have been written as c = 1*a and the 1 has been dropped as it would have had no mathematical impact on the calculation, this implies that they are perfectly dependent on each other in the same way that c = -1*a implies negative dependence. The calculation a + c would therefore need to be performed using perfect dependence. A third interpretation would be to consider that it is saying that c is a copy of a, they have the same uncertainty but their realisations are not necessarily related to each other but they have the same distribution shape. In the third scenario it would be sensible to make no assumptions about the dependencies between a and c and therefore Fréchet should be used. Knowledge of which of these scenarios is correct depends on the context of the script, something which Puffin is unable to make an assumption about by itself.

Another potential interpretation problem can occur because when creating code people naturally favour making their code readable. For example, the equation of motion for a damped harmonic oscillator can be given by

$$x''(t) + \frac{b}{m}x'(t) + \frac{k}{m}x(t) = 0 \tag{27}$$

where b is a damping constant, m is the mass of the oscillator and k is the spring constant. This equation can be solved analytically to find

$$x(t) = A \exp\left(\frac{-bt}{2m}\right) \cos\left(t\sqrt{\frac{k}{m} - \frac{b^2}{4m^2}} + \phi_0\right) \tag{28}$$

where A is a constant and ϕ_0 is the initial angle. In Figure 10 the equation has been coded in two different ways. In 10a the equation has been coded on a single line, as the equation is quite complicated it is likely that the programmer who is coding the equation would want to split it into multiple parts as has been done in 10b. There are no mathematical difference between the two approaches as they will lead to the same value. Puffin however would have to be careful around breaking up equations in such a way, strategies would be needed to ensure that breaking the lines up would not have a detrimental effect on how the code operated. Care would also need to be taken about the dependency tracking throughout the split equation, from example if there was uncertainty about the damping constant b then it would be difficult to assess the dependence between lines 4 and 5, especially since the cosine function is not monotonic. It may be better to use other techniques to solve the ODE (Equation 27) such as VSOPE or VNODE [53, 54, 55, 56, 57]. Making such a change would again require knowledge of what the calculation is and what exactly it is doing, something unlikely to be obvious from parsing the code.

6. Discussion

There are several reasons why Puffin may be unlikely to work as intended. The view that deterministic calculations can be translated to analogous computations for uncertainty quantification simply by replacing point values with uncertain structures like intervals and distributions ignores the issues of

1. computational burden,

16

2. repeated variables,

3. dependencies,

4. input specification,

5. conditionals,

6. ensembles, and

7. backcalculations.

This paper has described several strategies to mitigate the complexities arising from these issues. It is also likely to be the case that Puffin would not be able to introduce the most perfect uncertainty translations. Manually editing code or creating a new uncertainty aware script from scatch will likely outperform the automatic changes made by Puffin. This problem would not be unique to Puffin however, in general hand coding is always likely to outperform source-to-source translation [58].

Puffin is almost the exact opposite of an optimising compiler, which aims to minimise a programme's execution time and memory requirements. By replacing objects with uncertain equivalents both of these will almost certainly increase. For instance, if intervalising calculations increases the computational time 5- to 20-fold, clearly a simulation limited by computational time would need to be scaled back. Distributions or p-boxes would be still more burdensome. Of course, efficiency is not always a critical issue and this extra computational effort does pay for global uncertainty propagation and what computer scientists call automatic result verification [59]. Moreover, Puffin's implementation of modern uncertainty quantification could be more efficient and more comprehensive than simply embedding a deterministic computation inside a Monte Carlo shell with millions of replications.

It is unrealistic to expect that simply replacing computations involving integer and floating-point variables with analogous imprecise computations involving corresponding intervals, distributions and p-boxes will yield correct and useful results. The repeated variable problem and, more generally, the dependency problem, discussed in sections 4.6 and 5.2, can artifactually inflate the uncertainty of a complex computation. Even if the calculations are technically correct in the sense that they enclose the true uncertainty, the naive application of intrusive uncertainty quantification can sometimes yield results with massively inflated uncertainty that renders them practically useless. Having massive uncertainty is not the problem itself; the true uncertainty may actually be large. The problem is when it the uncertainty artifactually depends on the way the analysis was structured and does not reflect the features of the underlying computational problem.

The analysts who most need Puffin are those who have never heard of a p-box and who aren't sure what normal distributions or intervals are. Puffin offers multiple ways of specifying uncertainties in order to cater to the needs of such analysts. Puffin offers multiple ways of specifying uncertain inputs for user unfamiliar with the details that they require:

- Significant-digit intervals

- Measurement intervals (manufacturer or GUM conventions)

- English-language hedge words ('about', 'less than', etc.)

- Poisson model counts

- Moment or moment-range specifications

- Equivalent binomial count (k out of n confidence box)

- Single-sample confidence intervals

- Mean normal range (n & range normal distribution)

- Fermi strategies

- Distribution-free specification of p-boxes

It is also possible to run Puffin in a way that doesn't require any inputs, transforming assignments into significant digit intervals. Irrespective of how the uncertainties are expressed within Puffin, it is worth remembering that uncertain garbage in will lead to uncertain garbage out.

The general problem of uncertainty analysis is hard and it is difficult to create software that comprehensively solves all these problems and the development if Puffin is likely to be difficult. However, many practical problems are simpler than the most general problem and when this is the case it would be extremely useful to use a tool like Puffin which is able to handle uncertainty analysis intrusively within code. Puffin is intended to be open source and to be continuously co-developed by an interested community, adding in functionality and extending it as fashionable programming languages change and uncertainty quantification techniques develop.

Code Availability

Puffin is currently in development, and the current version can be found on GitHub[60]. We welcome suggestions and collaborators via GitHub.

References

[1] M. Shafto, M. Conroy, R. Doyle, E. Glaessgen, C. Kemp, J. LeMoigne, L. Wang, Modeling, Simulation, Information Technology and Processing Roadmap - Technology Area 11, Tech. rep., National Air and Space Administration, Washington, DC, USA (2012).

[2] S. Boschert, R. Rosen, Digital Twin—The Simulation Aspect, in: Mechatronic Futures, Springer International Publishing, Cham, 2016, pp. 59–74. doi:10.1007/978-3-319-32156-1.

[3] B. M. Adams, W. J. Bohnhoff, K. R. Dalbey, J. P. Eddy, M. S. Eldred, D. M. Gay, K. Haskell, P. D. Hough, L. P. Swiler, DAKOTA, A Multilevel Parallel Object-Oriented Framework for Design Optimization, Parameter Estimation, Uncertainty Quantification, and Sensitivity Analysis: Version 5.0 Reference Manual, Tech. rep., Sandia National Laboratories, Albuquerque, NM, United States (2010).

[4] E. Patelli, COSSAN: A Multidisciplinary Software Suite for Uncertainty Quantification and Risk Management, in: R. Ghanem, D. Higdon, H. Owhadi (Eds.), Handbook of Uncertainty Quantification, Springer International Publishing, Cham, 2015, pp. 1–69. doi:10.1007/978-3-319-11259-6_59-1.

[5] S. Marelli, B. Sudret, UQLab: A Framework for Uncertainty Quantification in Matlab, in: Vulnerability, Uncertainty, and Risk, American Society of Civil Engineers, Liverpool, UK, 2014, pp. 2554–2563. doi:10.1061/9780784413609.257.

[6] A. Olivier, D. G. Giovanis, B. Aakash, M. Chauhan, L. Vandanapu, M. D. Shields, UQpy: A general purpose Python package and development environment for uncertainty quantification, Journal of Computational Science 47 (2020) 101204. doi:10.1016/j.jocs.2020.101204.

[7] W. Oberkampf, Simulation informed decision making [Conference Presentation], in: Virtual Conference on Epistemic Uncertainty in Engineering, 2021, https://www.youtube.com/watch?v=i4L3fUpr59s.

[8] E. Paté-Cornell, On "Black Swans" and "Perfect Storms": Risk Analysis and Management When Statistics Are Not Enough, Risk Analysis 32 (11) (2012). `doi:10.1111/j.1539-6924.2011.01787.x`.

[9] R. P. Feynman, Appendix F - Personal Observations on Reliability of Shuttle, in: Report of the Presidential Commission on the Space Shuttle Challenger Accident, Vol. 2, US Government Printing Office, Washington, DC, USA, 1986, `https://history.nasa.gov/rogersrep/v2appf.htm`.

[10] Y. Amano, The Fukushima Daiichi Accident Report by the Director General, Tech. rep., International Atomic Energy Agency, Vienna, Austria (2015).

[11] M. S. Balch, R. Martin, S. Ferson, Satellite conjunction analysis and the false confidence theorem, Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences 475 (20180565) (2019). `doi:10.1098/rspa.2018.0565`.

[12] M. Beer, S. Ferson, V. Kreinovich, Imprecise probabilities in engineering analyses, Mechanical Systems and Signal Processing 37 (1-2) (2013) 4–29. `doi:10.1016/j.ymssp.2013.01.024`.

[13] D. A. Perez, H. Gietler, H. Zangl, Automatic Uncertainty Propagation Based on the Unscented Transform, in: 2020 IEEE International Instrumentation and Measurement Technology Conference (I2MTC), IEEE, Dubrovnik, Croatia, 2020, pp. 1–6. `doi:10.1109/I2MTC43012.2020.9129581`.

[14] E. Ilyushin, D. Namiot, On source-to-source compilers, International Journal of Open Information Technologies 4 (5.) (2016) 48–51, `http://www.scirp.org/journal/doi.aspx?DOI=10.4236/jsea.2013.64A005`.

[15] T. Parr, The Definitive ANTLR 4 Reference, The Pragmatic Bookshelf, Dallas, USA, 2012.

[16] Grammars written for ANTLR v4, `https://github.com/antlr/grammars-v4`.

[17] R. E. Moore, R. B. Kearfott, M. J. Cloud, Introduction to Interval Analysis, Vol. 110, Society for Industrial and Applied Mathematics, Philadelphia, USA, 2009.

[18] A. H.-S. Ang, W. S. Tang, Probability Concepts in Engineering : Emphasis on Applications in Civil & Environmental Engineering, 2nd Edition, John Wiley & Sons, Ltd, Hoboken, N.J., USA, 2007.

[19] S. Ferson, V. Kreinovich, L. Ginzburg, D. S. Myers, K. Sentz, Constructing Probability Boxes and Dempster-Shafer Structures, Tech. Rep. January, Sandia National Laboratories, Albuquerque, NM, United States (2003).

[20] M. S. Balch, Mathematical foundations for a theory of confidence structures, International Journal of Approximate Reasoning 53 (7) (2012) 1003–1019. `doi:10.1016/j.ijar.2012.05.006`.

[21] S. Ferson, J. O'Rawe, A. Antonenko, J. Siegrist, J. Mickley, C. C. Luhmann, K. Sentz, A. M. Finkel, Natural language of uncertainty: Numeric hedge words, International Journal of Approximate Reasoning 57 (2015) 19–39. `doi:10.1016/J.IJAR.2014.11.003`.

[22] P.-A. Jean, S. Harispe, S. Ranwez, P. Bellot, J. Montmain, Uncertainty detection in natural language: A probabilistic model, in: Proceedings of the 6th International Conference on Web Intelligence, Mining and Semantics, ACM, Nîmes France, 2016, pp. 1–10. `doi:10.1145/2912845.2912873`.

[23] S. Lefort, M.-J. Lesot, E. Zibetti, C. Tijus, M. Detyniecki, Interpretation of approximate numerical expressions: Computational model and empirical study, International Journal of Approximate Reasoning 82 (2017) 193–209. `doi:10.1016/j.ijar.2016.12.004`.

[24] W. F. Mascarenhas, Moore: Interval Arithmetic in C++20, in: G. A. Barreto, R. Coelho (Eds.), 37th Conference of the North American Fuzzy Information Processing Society, Vol. 831, Springer International Publishing, Fortaleza, Brazil, 2018, pp. 519–529. `doi:10.1007/978-3-319-95312-0_45`.

[25] Probability Bounds Analysis for Python 3, `https://pypi.org/project/pba/` or `https://github.com/Institute-for-Risk-and-Uncertainty/pba-for-python/`.

[26] Probability Bounds Analysis for MATLAB, `https://github.com/Institute-for-Risk-and-Uncertainty/pba-for-ma`

[27] Probability Bounds Analysis for R, `https://github.com/ScottFerson/pba.r`.

[28] HYRISK, `https://cran.r-project.org/web/packages/HYRISK/index.html`.

[29] Probability Bounds Analysis for Julia, `https://github.com/AnderGray/ProbabilityBoundsAnalysis.jl`.

[30] M. Haenggi, Meta Distributions–Part 1: Definition and Examples, IEEE Communications Letters (2021) 1–1`doi:10.1109/LCOMM.2021.3069662`.

[31] M. Haenggi, Meta Distributions–Part 2: Properties and Interpretations, IEEE Communications Letters (2021) 1–1`doi:10.1109/LCOMM.2021.3069681`.

[32] J. G. Dijkman, H. V. A. N. Haeringen, S. J. D. E. Lange, Fuzzy Numbers, Journal of Mathematical Analysis and Applicatins 92 (1983) 301–341.

[33] D. Dubois, H. Prade, Interval-valued Fuzzy Sets , Possibility Theory and Imprecise Probability, in: Proceedings of the Joint 4th Conference of the European Society for Fuzzy Logic and Technology and the 11th Rencontres Francophones Sur La Logique Floue et Ses Applications, Barcelona, Spain, 2005.

[34] M. S. Balch, New two-sided confidence intervals for binomial inference derived using Walley's imprecise posterior likelihood as a test statistic, International Journal of Approximate Reasoning 123 (2020) 77–98. `doi:10.1016/j.ijar.2020.05.005`.

[35] Y. Ben-Haim, Info-Gap Decision Theory: Decisions Under Severe Uncertainty, 2nd Edition, Academic Press, Oxford, UK, 2006.

[36] S. Ferson, R. B. Nelsen, J. Hajagos, D. J. Berleant, J. Zhang, W. T. Tucker, L. R. Ginzburg, W. L. Oberkampf, Dependence in probabilistic modeling, Dempster-Shafer theory, and probability bounds analysis, Tech. Rep. 19094, Sandia National Laboratories, Albuquerque, NM, USA (2004).

[37] S. Ferson, RAMAS Risk Calc 4.0 Software: Risk Assessment with Uncertain Numbers, Lewis Publishers, Boca Raton, Florida, USA, 2002, `https://books.google.co.uk/books?id=tKz7UZRs0CEC`.

[38] S. Ferson, M. Balch, K. Sentz, J. Siegrist, Computing with Confidence, in: Proceedings of the Eighth International Symposium on Imprecise Probability: Theory and Applications, Compiègne, France, 2013, `http://www.sipta.org/isipta13/proceedings/papers/s013.pdf`.

[39] A. Wimbush, N. Gray, S. Ferson, Singhing with Confidence: Visualising the Performance of Confidence Structures, arXiv:2106.04433 [stat]`http://arxiv.org/abs/2106.04433` (Jun. 2021). `arXiv:2106.04433`.

[40] L. D. Brown, T. T. Cai, A. Dasgupta, Interval Estimation for a Binomial Proportion, Statistical Science 16 (2) (2001) 101–117.

[41] C. J. Clopper, E. S. Pearson, The Use of Confidence or Fiducial Limits Illustrated in the Case of the Binomial, Biometrika 26 (4) (1934) 404–413.

[42] D. Hose, M. Hanss, A universal approach to imprecise probabilities in possibility theory, International Journal of Approximate Reasoning 133 (2021) 133–158. `doi:10.1016/j.ijar.2021.03.010`.

[43] V. Kreinovich, Decision Making Under Interval Uncertainty (and Beyond), in: P. Guo, W. Pedrycz (Eds.), Human-Centric Decision-Making Models for Social Sciences, Springer Berlin Heidelberg, Berlin, Heidelberg, 2014, pp. 163–193. `doi:10.1007/978-3-642-39307-5_8`.

[44] L. A. Zadeh, Fuzzy Logic, Computer 21 (4) (1988) 83–93. `doi:10.1109/2.53`.

[45] L. H. De Figueiredo, J. Stolfi, Affine arithmetic: Concepts and applications, Numerical Algorithms 37 (2004) 147–158.

[46] E. Goubault, S. Putot, A zonotopic framework for functional abstractions, Formal Methods in System Design 47 (3) (2016) 302–360. `doi:10.1007/s10703-015-0238-z`.

[47] A. Gray, M. De Angelis, S. Ferson, E. Patelli, What's Z-X, when Z = X+Y? dependency tracking in interval arithmetic with bivariate sets, in: 9th International Workshop on Reliable Engineering Computations, Virtual Conference, 2021, pp. 27–28.

[48] W. Krämer, Generalized Intervals and the Dependency Problem, Proceedings in Applied Mathematics and Mechanics 684 (2006) 683–684. `doi:10.1002/pamm.200610`.

[49] P. Embrechts, F. Lindskog, A. Mcneil, Modelling Dependence with Copulas and Applications to Risk Management, in: Handbook of Heavy Tailed Distributions in Finance, Elsevier, 2003, pp. 329–384. `doi:10.1016/B978-044450896-6.50010-8`.

[50] R. B. Nelsen, An Introduction to Copulas, 2nd Edition, Springer Series in Statistics, Springer, New York, New York, USA, 2006.

[51] H. Joe, Dependence Modeling with Copulas, Chapman & Hall/CRC, Boca Raton, Florida, USA, 2014.

[52] A. Gray, D. Hose, M. De Angelis, M. Hanss, S. Ferson, Dependent Possibilistic Arithmetic using Copulas, in: Proceedings of the Twelth International Symposium on Imprecise Probabilities: Theories and Applications, Vol. 147, Proceedings of Machine Learning Research, Granada, Spain (Virtual), 2021, pp. 173–183.

[53] N. S. Nedialkov, K. R. Jackson, G. F. Corliss, Validated solutions of initial value problems for ordinary differential equations, Appl. Math. Comput. (1999) 48.

[54] N. S. Nedialkov, K. R. Jackson, J. D. Pryce, An Effective High-Order Interval Method for Validating Existence and Uniqueness of the Solution of an IVP for an ODE, Reliable Computing 7 (6) (2001) 17.

[55] N. Nedialkov, Interval Tools for ODEs and DAEs, in: 12th GAMM - IMACS International Symposium on Scientific Computing, Computer Arithmetic and Validated Numerics (SCAN 2006), IEEE, Duisburg, Germany, 2006, pp. 4–4. `doi:10.1109/SCAN.2006.28`.

[56] Y. Lin, M. A. Stadtherr, Validated solutions of initial value problems for parametric ODEs, Applied Numerical Mathematics 57 (2007) 1145–1162. `doi:10.1016/j.apnum.2006.10.006`.

[57] J. A. Enszer, D. M. Andrei, M. A. Stadtherr, Probability bounds analysis for nonlinear population ecology models, Mathematical Biosciences 267 (2015) 97–108. `doi:10.1016/j.mbs.2015.06.012`.

[58] D. A. Plaisted, Source-to-Source Translation and Software Engineering, Journal of Software Engineering and Applications 06 (04) (2013) 30–40. `doi:10.4236/jsea.2013.64A005`.

[59] E. Adams, U. Kulisch, Scientific Computing with Automatic Result Verification, Academic Press, Boston, MA, USA, 1993.

[60] Puffin, https://github.com/ngg1995/Puffin.