

Invariance Explains Empirical Success of Many Intelligent Techniques

1st Olga Kosheleva
Department of Teacher Education
University of Texas at El Paso
El Paso, Texas 79968, USA
olgak@utep.edu

2nd Vladik Kreinovich
Department of Computer Science
University of Texas at El Paso
El Paso, Texas 79968, USA
vladik@utep.edu

Abstract—In many applications of intelligent computing, we need to choose an appropriate function – e.g., an appropriate re-scaling function, or an appropriate aggregation function. In applications of intelligent techniques, the problem of selecting an optimal function is usually too complex or too imprecise to be solved analytically, so the best functions are found empirically, by trying a large number of alternatives. In this paper, we show that in many such cases, the resulting empirical choice can be explained by natural invariance ideas. Example range from applications to building blocks of intelligent techniques – such as aggregation (including hierarchical aggregation) and averaging – to method-specific (polynomial fuzzy approach, pooling and averaging in deep learning) and domain-specific application, such as describing relative position of 2D and 3D objects, gauging segmentation quality, and perception of delay in public transportation.

Index Terms—intelligence techniques, re-scaling, aggregation, invariance

I. FORMULATION OF THE PROBLEM

Need to find dependencies. One of the main objectives of science is to predict the future state of the world, and one of the main objectives of engineering is to find out what can be done to make this future state better.

The state of the world is characterized by the values of relevant quantities. For example, to predict tomorrow's weather in a given area, we need to predict temperature, humidity, wind speed, wind direction, and other characteristics in different locations within the area of interest. To predict the future state of the world, we can use the current values of these and related quantities. So, to make a successful prediction, we need to know how exactly the future value of each quantity depends on these current values. In computational terms, what we need to know is an algorithm. In mathematical terms, what we need to know are functions describing such dependencies.

Similarly, our possible actions can also be characterized by different numerical parameters. To make the future state better, we need to know how the values of the future state depend

This work was supported in part by the National Science Foundation grants 1623190 (A Model of Change for Preparing a New Generation for Professional Practice in Computer Science), and HRD-1834620 and HRD-2034030 (CAHSI Includes), and by the AT&T Fellowship in Information Technology. It was also supported by the program of the development of the Scientific-Educational Mathematical Center of Volga Federal District No. 075-02-2020-1478, and by a grant from the Hungarian National Research, Development and Innovation Office (NRDI).

on these parameters. In such problems, we also need to know a function.

Need for intelligent techniques. In many physics situations and in many problems in other application areas, we know the desired function. For example, in celestial mechanics, we know exactly how to predict the future locations and velocities of celestial bodies based on their current locations and velocities – and based on our knowledge about the masses of these bodies.

However, in many other situations, we do not know the exact dependencies. In such situations, at best, we have some partial imprecise knowledge – knowledge that need to be handled by various intelligent techniques such as fuzzy, neural, etc.

Need for empirical selection of appropriate functions. These techniques provide many different options. Many of these options require selecting a function – e.g., selecting an activation function for a neural network. Usually in a problem, only some of possible options lead to a success. Which option is more successful depends on the situation. So, to be successful, we need to empirically select the best option – and thus, the best function.

What we do in this paper. In this paper, we summarize the results of empirically selecting best functions in different intelligent techniques. We show that in many such cases, the empirical success of the selected function can be explained by their invariance with respect to natural transformations.

These examples include techniques on all level of abstraction:

- techniques for basic sub-stages of intelligent techniques, e.g., for aggregation and averaging of data points;
- techniques specific for large segments of the corresponding intelligent techniques, e.g., the use of non-linear functions in Takagi-Sugeno-type fuzzy control and the use of pooling and averaging in deep learning; and
- intelligent techniques for analyzing specific objects: 1D time series (e.g., related to public transportation), 2D and 3D images (geographical and medical), etc.

From the mathematical viewpoint, our examples will fall into the following three categories:

- in some cases, we need to select a single function;

- in other cases, we need to select a family of functions; and
- in yet other cases, we need to select the best aggregation operation that combines several quantities and/or several functions.

II. INVARIANCE: A BRIEF REMINDER

Shift and scaling. We want to work with the actual values of physical quantities. What we actually work with are numerical values of these quantities. This is not just a linguistic distinction: the numerical value of a quantity depends not only on its actual values, it also depends on our selection of a measuring unit. The same height of 1.7 m gets a different numerical value 170 when described in centimeters. In general, if we switch from the original measuring unit to a new unit which is λ times smaller, then all numerical values of the corresponding quantity get multiplied by λ : $x \mapsto \lambda \cdot x$. This transformation is known as it *scaling*.

For many quantities like time and (macro-level) temperature, the numerical value also depends on our selection of the starting point. In general, if we select a new starting point which is x_0 units smaller than the previous one, then this value x_0 is added to all the numerical values: $x \mapsto x + x_0$. This transformation is known as *shift*.

Invariance. In many practical situations, there is no preferred measuring unit and no preferred starting point. In such situations, it makes sense to require that the corresponding dependencies not depend on these choices.

For example, if we are interested in the dependence $y = f(x)$ between two quantities x and y , and there is no preferred measuring unit for measuring x , then it is reasonable to require that the dependence look the same if we use a different measuring unit and, thus, different numerical values of the quantity x , namely, the values $X = \lambda \cdot x$.

Of course, we cannot simply require that the formula $y = f(x)$ remains the same: that would mean that $y = f(x) = f(X) = f(\lambda \cdot x)$ for all x and λ , which would imply that $f(x)$ is a constant. This may sound like a problem until we realize that, e.g., the formula $d = v \cdot t_0$ that described how far a body with velocity v can travel during time t_0 . This formula does not depend on what unit we use to describe velocity, but if we change a measuring unit for velocity, e.g., from km/h to miles per hour, we need to also appropriately change the measuring unit for distance.

Thus, we should require that for every $\lambda > 0$, there should be an appropriate transformation $y \mapsto Y$ for which $y = f(x)$ always implies $Y = f(X)$. This is what is usually meant by *invariance*. In this paper, we will show how invariance explains many empirical choices.

III. FUNCTIONS OF ONE VARIABLE: WHICH DEPENDENCIES ARE INVARIANT

General case. In general, if we change both the measuring unit and the starting point, we get a generic linear transformation $x \mapsto \lambda \cdot x + x_0$. So, in the ideal case, we should be looking for

dependencies $f(x)$ which are invariant with respect to both types of transformations.

In precise terms, such invariance means that for every $\lambda > 0$ and for every x_0 , there exists corresponding values $\mu > 0$ and y_0 depending on λ and x_0 , so that:

- every time we have $y = f(x)$,
- we also have $Y = f(X)$ for the same function f , where $X = \lambda \cdot x + x_0$ and $Y = \mu(\lambda, x_0) \cdot y + y_0(\lambda, x_0)$.

Substituting the expressions for X and Y into the formula $Y = f(X)$, we get

$$\mu(\lambda, x_0) \cdot y + y_0(\lambda, x_0) = f(\lambda \cdot x + x_0).$$

Here, $y = f(x)$, so we get

$$f(\lambda \cdot x + x_0) = \mu(\lambda, x_0) \cdot f(x) + y_0(\lambda, x_0).$$

Unfortunately, every measurable solution of this functional equation is a linear function (see, e.g., [29]). So, if we want to describe non-linear dependencies – and many real-life dependencies are non-linear – we cannot require invariance with respect to both scaling and shift, we can only require one type of invariance. Let us see what we can conclude if we make such requirements.

Possible cases. For each of the quantities x and y , we have two possible classes of transformations: scalings and shifts. Thus, we can have four possible cases:

- the case when a scaling of x leads to an appropriate scaling of y ;
- the case when a scaling of x leads to an appropriate shift of y ;
- the case when a shift of x leads to an appropriate scaling of y ; and
- the case when a shift of x leads to an appropriate shift of y .

Let us consider these four cases one by one. After that, we will show, on several examples, that invariance explains the empirical success of the corresponding intelligent techniques.

Scaling-scaling case. In this case, for every $\lambda > 0$, there exists a corresponding value $\mu > 0$ depending on λ so that:

- every time we have $y = f(x)$,
- we also have $Y = f(X)$ for the same function f , where $X = \lambda \cdot x$ and $Y = \mu(\lambda) \cdot y$.

Substituting the expressions for X and Y into the formula $Y = f(X)$, we get

$$\mu(\lambda) \cdot y = f(\lambda \cdot x).$$

Here, $y = f(x)$, so we get

$$f(\lambda \cdot x) = \mu(\lambda) \cdot f(x).$$

It is known (see, e.g., [1]) that every differentiable solution of this functional equation has the form

$$f(x) = A \cdot x^a$$

for some real numbers A and a (see also [29] for a detailed proof of this and related statements). Thus, in this case, we get the power law.

Comments.

- If we do not require differentiability, if we only require that the function $f(x)$ is measurable, then we can have different coefficient A for positive x and for negative x . A good example of such a not-everywhere-differentiable function is the Rectified Linear (ReLU) activation function $f(x) = \max(0, x)$ used in deep neural networks; see, e.g., [13].
- A similar formula can be obtained in the multi-dimensional case, when the desired quantity y depends on several variables x_1, \dots, x_n . In this case, we can select a new measuring unit for each of the n inputs, leading to $x_i \rightarrow \lambda_i \cdot x_i$. Thus, the requirement that the dependence not depend on the selection of measuring units means that for every tuples of the values $\lambda_1 > 0, \dots, \lambda_n > 0$, there exists a corresponding value $\mu > 0$ depending on $\lambda_1, \dots, \lambda_n$ so that:
 - every time we have $y = f(x_1, \dots, x_n)$,
 - we also have $Y = f(X_1, \dots, X_n)$ for the same function f , where $X_i = \lambda_i \cdot x_i$ and $Y = \mu(\lambda_1, \dots, \lambda_n) \cdot y$.
 Substituting the expressions for X and Y into the formula $Y = f(X)$, we get

$$\mu(\lambda_1, \dots, \lambda_n) \cdot y = f(\lambda_1 \cdot x_1, \dots, \lambda_n \cdot x_n).$$

Here, $y = f(x_1, \dots, x_n)$, so we get

$$f(\lambda_1 \cdot x_1, \dots, \lambda_n \cdot x_n) = \mu(\lambda_1, \dots, \lambda_n) \cdot f(x_1, \dots, x_n).$$

It is known (see, e.g., [1], [29]) that every differentiable solution of this functional equation has the form

$$f(x_1, \dots, x_n) = A \cdot x_1^{a_1} \cdot \dots \cdot x_n^{a_n}$$

for some real numbers A, a_1, \dots, a_n .

Scaling-shift case. In this case, for every $\lambda > 0$, there exists a corresponding value y_0 depending on λ so that:

- every time we have $y = f(x)$,
- we also have $Y = f(X)$ for the same function f , where $X = \lambda \cdot x$ and $Y = y + y_0$.

Substituting the expressions for X and Y into the formula $Y = f(X)$, we get

$$y + y_0(\lambda) = f(\lambda \cdot x).$$

Here, $y = f(x)$, so we get

$$f(\lambda \cdot x) = f(x) + y_0(\lambda).$$

It is known (see, e.g., [1], [29]) that every differentiable solution of this functional equation has the form

$$f(x) = A \cdot \ln(x) + a$$

for some real numbers A and a . Thus, in this case, we get a logarithmic dependence.

Shift-scaling case. In this case, for every x_0 , there exists a corresponding value $\mu > 0$ depending on x_0 so that:

- every time we have $y = f(x)$,
- we also have $Y = f(X)$ for the same function f , where $X = x + x_0$ and $Y = \mu(x_0) \cdot y$.

Substituting the expressions for X and Y into the formula $Y = f(X)$, we get

$$\mu(x_0) \cdot y = f(x + x_0).$$

Here, $y = f(x)$, so we get

$$f(x + x_0) = \mu(x_0) \cdot f(x).$$

It is known (see, e.g., [1], [29]) that every measurable solution of this functional equation has the form

$$f(x) = A \cdot \exp(a \cdot x)$$

for some real numbers A and a . Thus, in this case, we get an exponential dependence.

Shift-shift case. In this case, for every x_0 , there exists a corresponding value y_0 depending on x_0 so that:

- every time we have $y = f(x)$,
- we also have $Y = f(X)$ for the same function f , where $X = x + x_0$ and $Y = y + y_0(x_0)$.

Substituting the expressions for X and Y into the formula $Y = f(X)$, we get

$$y + y_0(x_0) = f(x + x_0).$$

Here, $y = f(x)$, so we get

$$f(x + x_0) = f(x) + y_0(x_0).$$

It is known (see, e.g., [1], [29]) that every measurable solution of this functional equation has the form

$$f(x) = A \cdot x + a$$

for some real numbers A and a . Thus, in this case, we get a linear dependence.

Comment. A similar formula can be obtained in the multi-dimensional case, when the desired quantity y depends on several variables x_1, \dots, x_n . In this case, we can select a starting point for each of the n inputs, leading to $x_i \mapsto x_i + x_{0i}$. Thus, the requirement that the dependence not depend on the selection of the starting point means that for every tuples of the values x_{01}, \dots, x_{0n} , there exists a corresponding value y_0 depending on x_{01}, \dots, x_{0n} so that:

- every time we have $y = f(x_1, \dots, x_n)$,
- we also have $Y = f(X_1, \dots, X_n)$ for the same function f , where $X_i = x_i + x_{0i}$ and $Y = y + y_0(x_{01}, \dots, x_{0n})$.

Substituting the expressions for X and Y into the formula $Y = f(X)$, we get

$$y + y_0(x_{01}, \dots, x_{0n}) = f(x_1 + x_{01}, \dots, x_n + x_{0n}).$$

Here, $y = f(x_1, \dots, x_n)$, so we get

$$f(x_1 + x_{01}, \dots, x_n + x_{0n}) = f(x_1, \dots, x_n) + y_0(x_{01}, \dots, x_{0n}).$$

It is known (see, e.g., [1], [29]) that every measurable solution of this functional equation is a linear function

$$f(x_1, \dots, x_n) = A + a_1 \cdot x_1 + \dots + a_n \cdot x_n,$$

for some real numbers A, a_1, \dots, a_n .

Applications of these results. Intelligent technique process data. This can be 1D data – e.g., several values of measuring the same quantity, or it can be 2D (or even 3D) data corresponding to images. Let us show that the above invariance result can be used to explain empirical successes of intelligent techniques in processing both types of data.

Invariance explains empirical successes of intelligent techniques for processing 1D data. There are many effective techniques for data processing, when we have the results x_1, \dots, x_n of measuring or estimating several quantities, and we need to estimate the values of related quantities $y = f(x_1, \dots, x_n)$ – e.g., predicting tomorrow's weather based on today's meteorological data (and on the historical meteorological data). Somewhat surprisingly, the existing techniques are not as good in a seemingly much simpler 1D problem: when we have several measurements and/or estimates x_1, \dots, x_n of the same quantity x , and we would like to “average” them, i.e., to combine them into a single estimate.

In practice, we usually assume that the estimation errors are independent and normally distributed. The normal distribution assumption is justified by the Central Limit Theorem (see, e.g., [30]), according to which the joint effect of many small independent factors – which is usually the case – is close to normal. Under this assumption, if we do not have any information about which estimate is more accurate, then it is natural to assume that all estimates are equally accurate. In this case, the optimal resulting estimate is the arithmetic average [30]:

$$\bar{x} = \frac{1}{n} \cdot \sum_{i=1}^n x_i.$$

In other cases, we know the accuracy of each estimate, i.e., in precise terms, we know the corresponding standard deviations σ_i . In such cases, the optimal estimate has the form

$$\bar{x} = \sum_{i=1}^n w_i \cdot x_i$$

for some weights $w_i \geq 0$ for which

$$\sum_{i=1}^n w_i = 1.$$

The optimal weights depend on the standard deviations σ_i :

$$w_i = \frac{\sigma_i^{-2}}{\sum_{j=1}^n \sigma_j^{-2}}.$$

The problem is that for many quantities, we have many different scales, and the average depends on what scale we use. For example, we can describe the strength of an earthquake

by the released energy, or when we can describe it by the logarithm of this energy. So, instead of averaging the original values x_i , we can average the transformed values $z_i = g(x_i)$, for some non-linear function $g(x)$, resulting in

$$\bar{z} = \sum_{i=1}^n w_i \cdot z_i = \sum_{i=1}^n w_i \cdot g(x_i),$$

and then re-scale back into the original scale, resulting in

$$\bar{x} = g^{-1} \left(\sum_{i=1}^n w_i \cdot g(x_i) \right),$$

where $g^{-1}(z)$ denote the inverse function.

Which function $g(x)$ should be use to get the most accurate results? It is reasonable to require that the transformation $g(x)$ is invariant. Depending on which type of invariance we assume, we get either power law, or logarithmic dependence, or the exponential dependence, or the linear function. For linear functions, the above expression leads back to arithmetic average. For other cases, we get different averaging operations. Namely, power-law $g(x)$ leads to

$$\bar{x} = \left(\sum_{i=1}^n w_i \cdot x_i^a \right)^{1/a},$$

logarithmic dependence $g(x)$ leads to

$$\bar{x} = \prod_{i=1}^n x_i^{w_i},$$

and the exponential dependence leads to

$$\bar{x} = \frac{1}{a} \cdot \ln \left(\sum_{i=1}^n w_i \cdot \exp(a \cdot x_i) \right).$$

In the limits $a \rightarrow \infty$ and $a \rightarrow -\infty$, the power-law formula leads to

$$\bar{x} = \max(x_1, \dots, x_n)$$

and to

$$\bar{x} = \min(x_1, \dots, x_n).$$

Interestingly, these are exactly the empirically observed averaging operations; see, e.g., [2], [3]. Thus, in this case, the empirical effectiveness is also explained by invariance.

Comment. In some cases, the most effective averaging operations are more complex. One example of such an averaging operation is so-called Lehmer means [2], [3], [7]

$$\bar{x} = \frac{\frac{1}{n} \cdot \sum_{i=1}^n w_i \cdot x_i^a}{\frac{1}{n} \sum_{i=1}^n w_i \cdot x_i^{a-1}}.$$

This operation can also be explained by scale-invariance: namely, it is the result of applying a scale-invariant function of two variables

$$y = f(y_1, y_2) = y_1^a \cdot y_2^{-(a-1)}$$

to averaging operations corresponding to scale-invariant functions $g_1(x) = x^a$ and $g_2(x) = x^{a-1}$:

$$y_1 = \left(\sum_{i=1}^n w_i \cdot x_i^a \right)^{1/a} \quad \text{and} \quad y_2 = \left(\sum_{i=1}^n w_i \cdot x_i^{a-1} \right)^{1/(a-1)}.$$

Other empirically successful operations are the following operations [2], [3], [34]:

$$\bar{x} = p \cdot \frac{1}{n} \cdot \sum_{i=1}^n x_i + (1-p) \cdot \min(x_1, \dots, x_n)$$

and

$$\bar{x} = p \cdot \frac{1}{n} \cdot \sum_{i=1}^n x_i + (1-p) \cdot \max(x_1, \dots, x_n)$$

for some $p \in [0, 1]$. These operations are obtained if we apply a shift-invariant combination operation – i.e., linear combination – to two invariant averaging operations: average and min (or max).

Invariance explains empirical successes of intelligent techniques for processing 2D and 3D data. Current image processing techniques have achieved great performance in image processing – e.g., in detecting objects in images, in describing different types of well-defined relationship between these objects; see, e.g., [13]. The current methods are, however, not as successful in describing intuitive, imprecise relationship between objects. We humans are accustomed to describe relative position of objects in imprecise terms: close by, far away, somewhat to the East, etc. We use these descriptions to make decisions. It is therefore necessary to be able, given a scene, to generate an appropriate description of relative positions of different objects in such understandable natural-language terms.

For this task, one of the most efficient methods is a Histogram of Forces method, in which:

- for each direction,
- we integrate the force $F(r)$ over all the lines parallel to this direction.

Here r is the shortest distance between points from these two objects that happen to be on this particular line; see, e.g., [5], [8]–[10], [21]. The effectiveness of this method depends on the choice of the function $F(r)$.

Empirically, it turned out that the most effective functions are the power laws $F(r) = A \cdot r^a$. The most widely used cases are $a = 0$ (constant force) and $a = -2$ (when the expression is the same as for the gravitational force between two bodies), but other values of a have also shown to be effective in some cases, e.g., the value $a = 2$ [9], [10].

The power laws are exactly the functions which are scale-scale invariant. Since, as we have mentioned, scale-scale invariance is a natural property, this explains the empirical success of these functions.

Comment. In some cases, other functions $F(r)$ turned out to be the most effective, e.g., the functions $F(r) = \min(C, r^{-2})$

and $F(r) = r_0^2 / (r_0^2 + r^2)$; see, e.g., [5], [21]. In the following text, we will show that the general invariance ideas explain the effectiveness of such functions as well.

IV. WHICH FAMILIES OF FUNCTIONS ARE INVARIANT

Why we need families of functions. In the previous section, we were looking for a single function that would be most efficient in solving several problems. Often, in different applications, different functions are more effective. In this case, instead of looking for a *single* function, it makes sense to look for a finite-parametric *family* of functions, so that we would be able to adjust the values of the corresponding parameters for each individual case.

Let us analyze which families are invariant.

What families we will consider. The simplest families are linear combinations of known functions, i.e., families of the type

$$C_1 \cdot f_1(x) + \dots + C_k \cdot f_k(x),$$

where the functions $f_1(x), \dots, f_k(x)$ are fixed, and C_1, \dots, C_k are the parameters that can be adjusted.

Which families are scale-invariant and shift-invariant. In the case of a single function, we could not require invariance with respect to both changing the measuring unit and changing the starting point – if we require both, we only get linear functions.

Interestingly, for families, it *is* possible to require both invariance. It turns out (see, e.g., [24]) that the only family of functions which is both scale-invariant and shift-invariant is the family of polynomials of a given order, i.e., the family of all the functions of the type

$$f(x) = C_1 + C_2 \cdot x + C_3 \cdot x^2 + \dots + C_k \cdot x^{k-1}.$$

Comments.

- A similar results holds if we consider scale-invariant and shift-invariant families of functions of several variables: all elements of such functions are polynomials.
- It is also possible to require only shift-invariance or only scale-invariance; see [24]. If we require only shift-invariance, then all functions from the invariant family are linear combinations of the expressions $x^m \cdot \exp(a \cdot x)$, where m is non-negative integer and a is, in general, a complex number. If we require only scale-invariance, then all functions from the invariant family are linear combinations of the expressions $(\ln(x))^m \cdot x^a$, where also m is non-negative integer and a is, in general, a complex number.

Applications of this result. For this result, we will describe two applications: a general application to intelligent systems and a more specific application to human-oriented systems.

General application to intelligent systems. As we have mentioned, in many situations of prediction and control, we

do not have exact knowledge of the system that we want to control, but we have knowledge of experts who have been successfully predicting and/or successfully controlling such systems. The problem is that experts formulate their knowledge in imprecise (“fuzzy”) terms, by using imprecise words from natural language. To reformulate such a knowledge in precise computer-understandable terms, Zadeh invented fuzzy techniques; see, e.g., [4], [15], [22], [26], [27], [33].

One of the most successful ways to use this technique in control is to use Takagi-Sugeno approach, where, to describe how experts predict or control based on the inputs x_1, \dots, x_n , we look for expert rules of the type

$$\text{if } A_1(x_1), \dots, \text{ and } A_n(x_n), \text{ then } y = f(x_1, \dots, x_n),$$

when A_i are fuzzy properties and $f(x_1, \dots, x_n)$ is a linear function.

The pioneering paper [31] showed that more effective prediction and control can be obtained if we allow polynomial functions $f(x_1, \dots, x_n)$. This approach is known as *polynomial fuzzy approach*. Since invariance is a natural property, and all the functions from scale-invariant and shift-invariant families are polynomials, this explains the empirical success of polynomial fuzzy approach.

Application to human-oriented systems. Intelligent control techniques are largely used in situations when the objective is clear. For human-oriented system, an additional challenge is that for such systems, the goal is subjective, it deals with perceptions and not with objective characteristics. To design such systems, we need to be able to predict such perceptions, i.e., we need to know how a human will react to different situations.

As a case study, let us consider the planning of public transportation – an important feature of big cities and of their smart-city plans. The goal of planning is to make the public transportation as convenient to people as possible. The main source of inconvenience is waiting time. Part of the waiting time is caused by the fact that the buses (and other means of public transportation) go by schedule, so a passenger has to wait for the next scheduled bus. This is a known inconvenience, people get adjusted to it – by taking the bus schedule into account when planning their trips. A much more serious inconvenience occurs when the buses are behind schedule. Such situations are unpredictable, they interfere with people’s plans and cause a lot of frustration. According to [32], there are several levels of such frustration, corresponding to 10, 20, and 60 minutes: delay below 10 min are perceived as negligible, delays between 10 and 30 min are perceived as short, delays between 30 and 60 minutes are perceived as long, and delays longer than 60 minutes are perceived as severe.

According to decision theory (see, e.g., [11], [12], [16], [19], [23], [25], [28]), people’s attitudes can be described by a function called utility – the larger the utility the more beneficial the situation. Utility is defined modulo possible linear transformation $u \mapsto \lambda \cdot u + u_0$ for some $\lambda > 0$ and u_0 .

Very small changes in a situation lead to very small, barely noticeable changes in utility. Let us denote by Δu the smallest noticeable change in utility. In this case, if we take, as a starting point for measuring utility, the value corresponding to no delay, the first noticeable change 10 min correspond to utility $-\Delta u$, the second (30 min) to $-2\Delta u$, and the third one (60 min) to $-3\Delta u$.

In general, how can we describe the dependence $t = f(u)$ of delay time t on utility u ? This dependence may be different for different people, so it makes sense to select not a single function $t = f(u)$, but a whole family of functions. Since utility is defined modulo scalings and shifts, it is reasonable to require that this family should be scale-invariant and shift-invariant. Thus, it must be a family of polynomials.

The simplest polynomials are linear functions, but a linear dependence would mean that a change from 5 min delay to 15 min delay is as painful as a change from 60 min delay to 70 min delay. In reality, if we have already waited for the bus for the whole hours, additional 10 minutes are not as painful as in the first case, when the delay time triples. To capture this difference – which is not reflected in the linear dependence – we need to go beyond linear functions. The simplest family of polynomials that includes non-linear functions is the family of all quadratic polynomials. And indeed, quadratic functions perfectly explain the above empirical dependence: if we select parameters C_i of the quadratic function $t(u) = C_1 + C_2 \cdot u + C_3 \cdot u^2$ to make sure that $t(0) = 0$, $t(-\Delta u) = 10$, and $t(-2\Delta u) = 30$, we will get

$$C_1 = 0, \quad C_2 = -\frac{5}{\Delta u}, \quad C_3 = \frac{5}{(\Delta u)^2}.$$

For these values, we get $t(-3\Delta u) = 60$, exactly what the empirical data shows.

So, in this case, invariance also explains an empirical dependence.

Comment. Interestingly, the dependence of amount of money on the utility is also quadratic – see, e.g., [14] – and for the same reason: a change from \$5 to \$15 causes much more positive feelings than a change from \$60 to \$70.

V. WHICH AGGREGATION OPERATIONS ARE INVARIANT

What is an aggregation operation. In many practical situations, we need to combine (aggregate) two or more values. For example, when we know the masses m_1 and m_2 of two objects, their total mass m is equal to the sum of their masses: $m = m_1 + m_2$. When we have two independent random variables with known standard deviations σ_1 and σ_2 , then the standard deviation σ of their sum is equal to $\sqrt{\sigma_1^2 + \sigma_2^2}$. The function that transforms the original values a_1 and a_2 into a new value is known as an *aggregation operation*. We will denote such operations by $a_1 * a_2$.

The combination result should not depend on the order, so it is usually assumed that this operation should be commutative: $a_1 * a_2 = a_2 * a_1$. For the same reason, often, associativity is required. In most practical situations, it is reasonable to require

monotonicity: if $a_1 \leq a'_1$ and $a_2 \leq a'_2$, then we should have $a_1 * a_2 \leq a'_1 * a'_2$.

Aggregation operations and averaging operations. Once we have an aggregation operation $a_1 * a_2$, we can define the corresponding averaging operation: it transforms the values a_1, \dots, a_n into the value a for which $a_1 * \dots * a_n = a * \dots * a$. For example, if we start with the sum $a_1 * a_2 = a_1 + a_2$, then we get arithmetic average

$$\frac{a_1 + \dots + a_n}{n}.$$

If we start with the product $a_1 * a_2 = a_1 \cdot a_2$ then we get geometric average $\sqrt[n]{a_1 \cdot \dots \cdot a_n}$, etc.

Which aggregation operations are invariant. Let us first consider scale-invariance, i.e., the property that if $a = a_1 * a_2$, then, for every $\lambda > 0$, we should have $A = A_1 * A_2$, where $A = \lambda \cdot a$, $A_1 = \lambda \cdot a_1$, and $A_2 = \lambda \cdot a_2$. It is known (see, e.g., [17]) scale-invariant, commutative, associative, continuous, and monotonic aggregation operations are: $a_1 * a_2 = (a_1^v + a_2^v)^{1/v}$, its limit cases $a_1 * a_2 = \max(a_1, a_2)$ and $a_1 * a_2 = \min(a_1, a_2)$ corresponding to $v \rightarrow \infty$ and $v \rightarrow -\infty$, and the trivial operation for which $a_1 * a_2 = 0$ for all a_1 and a_2 .

If we also require shift-invariance, i.e., require that $a = a_1 * a_2$ imply that $A = A_1 * A_2$, where $A = a + a_0$, $A_1 = a_1 + a_0$, and $A_2 = a_2 + a_0$, then the only remaining non-trivial aggregation operations are min and max. If instead we require a weaker version of shift-invariance, where $a = a_1 * a_1$ imply that $A = A_1 * A_2$ for $A_1 = a_1 + a_0$, $A_2 = a_2 + a_0$, and $A = a + a'_0(a_0)$ for some value a'_0 depending on a_0 , then we also get addition – which, as we have mentioned, corresponds to the arithmetic average.

Applications. As promised, we will show that the above results explain effectiveness of empirical intelligent techniques on all levels of abstraction:

- on the level of general techniques – in this case, it will be deep learning,
- on the level of specific applications – in this case, it will be applications to imaging, and
- on the level of building blocks – in this case, it will be aggregation.

Applications to deep learning. One of the main ideas behind deep learning – as compared to few-layers traditional neural networks – is that we have more layers and, correspondingly, fewer neurons in each layer. In the traditional neural network, when we had a large number of neurons in each layer – in particular, in the input layer – we could allocated, to each input value, a corresponding input neuron. In deep learning, in situations when we process a large amount of data – e.g., pixels forming an image – there is much fewer neurons in the input layer than there are inputs. So, we need to combine several inputs into a single value. This process is known as *pooling*. Since it is reasonable to require that this operation be scale-invariant and – at least weakly – shift-invariant, it is reasonable

to use max-pooling, min-pooling, or average pooling. These are indeed three most widely used pooling operations in deep learning [13]. Thus, the empirical success of these operations can also be explained by invariance [17].

Another problem of machine learning in general – and of deep learning in particular – is that its results are not always reliable. In general, a natural way to increase the reliability is to duplicate efforts, i.e., to have several similar devices working in parallel, and then to somehow average their results. This is, e.g., how we get the most accurate time – by having three or more super-precise clocks working in parallel and by averaging their results. Similarly, to increase reliability of a neural network, we simultaneously train several networks – i.e., in effect, subnetworks of the overall network, and then average the results. Thus, it makes sense to use averaging corresponding to invariant aggregations, e.g, arithmetic average – which is also one of the two most empirically effective averaging methods [13], [17]. (The other empirically effective method is geometric average that corresponds to $v \rightarrow 0$.)

First application to image processing. Now it is time to go to the image processing example. As we have mentioned earlier, scale-invariance leads to the power law $F(r) = A \cdot r^a$. In some cases, some values a are better, in other cases, other values of a are better. It is therefore reasonable to try to aggregate functions corresponding to different values a , so that the resulting function combine the advantages of both aggregated expressions. By applying a scale-invariant aggregation, we get either an expression $((A_1 \cdot r^{a_1})^v + (A_2 \cdot r^{a_2})^v)^{1/v}$, or $\max(A_1 \cdot r^{a_1}, A_2 \cdot r^{a_2})$, or $\min(A_1 \cdot r^{a_1}, A_2 \cdot r^{a_2})$. In particular, if you apply the simplest of such combinations – min – to the most successful cases $a_1 = 0$ and $a_2 = -2$, we get one of the hybrid force formulas $F(r) = \min(C, r^{-2}, C)$ that was empirically shown to be effective. For $v = -1$, $A_1 = 1$, $A_2 = r_0^2$, $a_1 = 0$, and $a_2 = -2$, we get another empirically successful formula $F(r) = r_0^2 / (r_0^2 + r^2)$. Thus, invariance explains these empirical successes as well.

Second application to image processing: gauging quality of skin lesion segmentation. Not only we need to make intelligent techniques more reliable, we also need to be able to gauge how reliable they are. A study [18] provides a new method for this gauging, which the authors illustrate on the case study of skin lesion segmentation. This new method combines taking the arithmetic average – of several images provided by different subnetworks, and taking min – namely, the minimum of the distances from each pixel to different points on the boundaries between the segments. These are exactly the scale-invariant and (weakly) shift-invariant aggregation operation. Thus, invariance explains the empirical success of (at least this part of) the new method.

Application to building blocks of intelligent techniques: hierarchical aggregation. Aggregation does not have to be performed in one step. For example, to get an average temperature on campus, it makes sense first to aggregate temperature values within each room – if the room has several sensors, then

combine these values to get an average of rooms, and then aggregate the building averages to get the overall campus average. The study of different aggregation techniques described in [20] showed that – in line with the above-mentioned result – the best results emerge when on each aggregation stage, we use min, max, or arithmetic average. Thus, this empirical result is also explained by invariance.

Comment. The paper [20] has an additional empirical observation: that on the first aggregation levels, it is advantageous to use min and max, while on the following levels, arithmetic average works better. This fact can be explained by the fact that with some small probability, sensors malfunction, and produce readings which are much larger or much smaller than the actual temperature. So:

- On the earlier aggregation stages, when we combine the readings of a small number of sensors, the probability that one of the readings is an outlier is still small. So, with probability close to 1, the max and min of these readings reflect the actual highest and lowest temperature in the room.
- On the other hand, on the later aggregation stages, we combine, in effect, a large number of readings. In this cases, there is a high probability that at least one of the combined values is an outlier. Thus, if we simply use max or min to aggregate, with high probability, we will get this outlier – and not the desired highest or lowest temperature on campus. Since we cannot use max or min, the only remaining option is to compute the arithmetic average.

ACKNOWLEDGMENT

The authors are greatly thankful to Janusz Kacprzyk for his encouragement, and to all the participants of the 2022 World Congress on Computational Intelligence WCCI'2022 (Padua, Italy, July 18–23, 2022) for fruitful discussions.

REFERENCES

- [1] J. Aczél and J. Dhombres, *Functional Equations in Several Variables*, Cambridge University Press, 2008.
- [2] J. G. Bazan, S. Bazan-Socha, U. Bentkowska, W. Galka, M. Mrukowicz, and L. Zaręba, “Comparison of aggregation classes in ensemble classifiers for high dimensional datasets”, *Proceedings of the 2022 World Congress on Computational Intelligence WCCI'2022*, Padua, Italy, July 18–23, 2022.
- [3] G. Beliakov, H. Bustince, and T. Calvo, *A Practical Guide to Averaging Functions*, Springer, Cham, Switzerland, 2016.
- [4] R. Belohlavek, J. W. Dauben, and G. J. Klir, *Fuzzy Logic and Mathematics: A Historical Perspective*, Oxford University Press, New York, 2017.
- [5] A. R. Buck, D. T. Anderson, J. M. Keller, R. H. Luke III, and G. Scott, “A comparison of relative position descriptors for 3D objects”, *Proceedings of the 2022 World Congress on Computational Intelligence WCCI'2022*, Padua, Italy, July 18–23, 2022.
- [6] A. R. Buck, J. M. Keller, and M. Skubic, “A memetic algorithm for matching spatial configurations with the histograms of forces”, *IEEE Transactions on Evolutionary Computation*, 2013, Vol. 17, No. 4, pp. 588–604.
- [7] P. S. Bullen, D. S. Mitrinović, and P. M. Vasić, *Means and Their Inequalities*, Reidel, Dordrecht, 1988.
- [8] M. Clément, C. Kurtz, and L. Wendling, “Bags of spatial relations and shapes features for structural object description”, *Proceedings of the 23rd IEEE International Conference on Pattern Recognition ICPR'2016*, Cancún, Mexico, December 4–8, 2016, pp. 1995–2000.
- [9] M. Clément, C. Kurtz, and L. Wendling, “Learning spatial relations and shapes for structural object description and scene recognition”, *Pattern Recognition*, 2018, Vol. 84, pp. 197–210.
- [10] R. Deléarde, C. Kurtz, P. Dejean, and L. Wendling, “Force Banner for the recognition of spatial relations”, *Proceedings of the 25th International Conference on Pattern Recognition ICPR'2020*, Milan, Italy, January 10–15, 2021, pp. 6065–6072.
- [11] P. C. Fishburn, *Utility Theory for Decision Making*, John Wiley & Sons Inc., New York, 1969.
- [12] P. C. Fishburn, *Nonlinear Preference and Utility Theory*, The John Hopkins Press, Baltimore, Maryland, 1988.
- [13] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*, MIT Press, Cambridge, Massachusetts, 2016.
- [14] D. Kahneman, *Thinking, Fast and Slow*, Farrar, Straus, and Giroux, New York, 2011.
- [15] G. Klir and B. Yuan, *Fuzzy Sets and Fuzzy Logic*, Prentice Hall, Upper Saddle River, New Jersey, 1995.
- [16] V. Kreinovich, “Decision making under interval uncertainty (and beyond)”, In: P. Guo and W. Pedrycz (eds.), *Human-Centric Decision-Making Models for Social Sciences*, Springer Verlag, 2014, pp. 163–193.
- [17] V. Kreinovich and O. Kosheleva, “Optimization under uncertainty explains empirical success of deep learning heuristics”, In: P. Pardalos, V. Rasskazova, and M. N. Vrahatis (eds.), *Black Box Optimization, Machine Learning and No-Free Lunch Theorems*, Springer, Cham, Switzerland, 2021, pp. 195–220.
- [18] Q. Lin, X. Chen, C. Chen, and J. M. Garibaldi, “Quality quantification in deep convolutional neural networks for skin lesion segmentation using fuzzy uncertainty measurement”, *Proceedings of the 2022 World Congress on Computational Intelligence WCCI'2022*, Padua, Italy, July 18–23, 2022.
- [19] R. D. Luce and R. Raiffa, *Games and Decisions: Introduction and Critical Survey*, Dover, New York, 1989.
- [20] L. Magdalena, L. Garmendia, D. Gómez, and J. Montero, “Hierarchical computable aggregations”, *Proceedings of the 2022 World Congress on Computational Intelligence WCCI'2022*, Padua, Italy, July 18–23, 2022.
- [21] P. Matsakis, *Relations spatiales structurelles et interprétation d'images*, Ph.D. dissertation, IRIT, Université Paul Sabatier, Toulouse, France, 1998.
- [22] J. M. Mendel, *Uncertain Rule-Based Fuzzy Systems: Introduction and New Directions*, Springer, Cham, Switzerland, 2017.
- [23] H. T. Nguyen, O. Kosheleva, and V. Kreinovich, “Decision making beyond Arrow's ‘impossibility theorem’, with the analysis of effects of collusion and mutual attraction”, *International Journal of Intelligent Systems*, 2009, Vol. 24, No. 1, pp. 27–47.
- [24] H. T. Nguyen and V. Kreinovich, *Applications of Continuous Mathematics to Computer Science*, Kluwer, Dordrecht, 1997.
- [25] H. T. Nguyen, V. Kreinovich, B. Wu, and G. Xiang, *Computing Statistics under Interval and Fuzzy Uncertainty*, Springer Verlag, Berlin, Heidelberg, 2012.
- [26] H. T. Nguyen, C. L. Walker, and E. A. Walker, *A First Course in Fuzzy Logic*, Chapman and Hall/CRC, Boca Raton, Florida, 2019.
- [27] V. Novák, I. Perfilieva, and J. Močkoř, *Mathematical Principles of Fuzzy Logic*, Kluwer, Boston, Dordrecht, 1999.
- [28] H. Raiffa, *Decision Analysis*, McGraw-Hill, Columbus, Ohio, 1997.
- [29] E. D. Rodriguez Velasquez, V. Kreinovich, and O. Kosheleva, “Invariance-based approach: general methods and pavement engineering case study”, *International Journal of General Systems*, 2021, Vol. 50, No. 6, pp. 672–702.
- [30] D. J. Sheskin, *Handbook of Parametric and Nonparametric Statistical Procedures*, Chapman and Hall/CRC, Boca Raton, Florida, 2011.
- [31] K. Tanaka, H. Yoshida, H. Ohtake, and H. O. Wang, “A sum-of-squares approach to modeling and control of nonlinear dynamical systems with polynomial fuzzy systems”, *IEEE Transactions on Fuzzy Systems*, 2009, Vol. 17, No. 4, pp. 911–922.
- [32] N. D. T. Tran, C. K. Leung, T. Turner, S. T. Wu, N. Karimbaeva, J. Kim, and A. Cuzzocrea, “Transportation analytics with fuzzy logic and regression”, *Proceedings of the 2022 World Congress on Computational Intelligence WCCI'2022*, Padua, Italy, July 18–23, 2022.
- [33] L. A. Zadeh, “Fuzzy sets”, *Information and Control*, 1965, Vol. 8, pp. 338–353.
- [34] H. J. Zimmermann and P. Zysno, “Decisions and evaluations by hierarchical aggregation of Information”, *Fuzzy Sets and Systems*, 1983, Vol. 10, pp. 243–260.