Tables as Powerful Representational Tools

Dirk Schlimm(✉)

McGill University, Montreal, QC H3A 2T7, Canada
dirk.schlimm@mcgill.ca

Abstract. Tables are widely used for storing, retrieving, communicating, and processing information, but in the literature on the study of representations they are still somewhat neglected. The strong structural constraints on tables allow for a clear identification of their characteristic features and the roles these play in the use of tables as representational and cognitive tools. After introducing syntactic, spatial, and semantic features of tables, we give an account of how these affect our perception and cognition on the basis of fundamental principles of Gestalt psychology. Next are discussed the ways in which these features of tables support their uses in providing a global access to information, retrieving information, and visualizing relational structure and patterns. The latter is particularly important, because it shows how tables can contribute to the generation of new knowledge. In addition, tables also provide efficient means for manipulating information in general and in structured notations. In sum, tables are powerful and efficient representational tools.

Keywords: Tables · Diagrams · Cognition · Gestalt psychology

1 Introduction

In the wake of the groundbreaking analysis of the use of diagrams by Larkin and Simon [15], the computational efficacy of a representational system for solving particular tasks has become an important criterion for the assessment of such systems. According to Giardino, this raises the question of how this efficacy 'happens to emerge from the interaction between more spontaneous abilities and the production of cultural artifacts' [6, p. 81]. While by far the most work in this area has been done on diagrams, understood more generally, in the present paper a somewhat neglected kind of representation is discussed, namely *tables*. As we shall see below, tables are fairly constrained representations, which explains perhaps why they have not received much attention to so far. If they are mentioned at all, tables are mainly discussed together with other kinds of representations, — e. g., in [7] they are considered together with graphs and illustrations, in [11] and [21] as instances of a much broader category of diagrams, — so that their specific features have often been glossed over. A noteworthy exception are the analyses of relational tables by Shimojima and his colleagues [19,20,26].

The main questions addressed in this paper are: what are the specific features of tables and how do these contribute to the particular uses of tables?

V. Giardino et al. (Eds.): Diagrams 2022, LNAI 13462, pp. 185–201, 2022.
https://doi.org/10.1007/978-3-031-15146-0_15

The structural constraints of tables allow for an explicit identification of their syntactic, spatial, and semantic features and, in turn, of the roles these play in the use of tables as representational and cognitive tools.

In the discussion we attempt to navigate between the Scylla of studying very specific kinds of systems and tasks, which can lead to concrete insights, but raises the question of their generalizability, and the Charybdis of trying to cover a wide range of different systems and thereby running the risk of leading to a high-level analysis that has only little to say about specific cases. In the following, we begin by giving an overview of what we consider tables to be, of their main syntactic, spatial, and semantic features and of the effects of these features on perception and cognition (Sect. 2). We then turn to studying the relation of these features to the, so to speak, *passive* uses of tables: for the general presentation of information, information retrieval, and the visualization of relational structure and patterns (Sects. 3–5). The latter is particularly important, because it shows how tables can contribute to the detection of new patterns and generation of new knowledge. In Sects. 6–8, *active* interactions with tables are discussed, namely operations on tables, tabular manipulations of structured notations, and operations on infinite tables. Together, these discussions aim to shed some light on the power of tables as representational tools and on how this comes about.

2 Tables and Their Features

Definition. Tables are defined here in terms of syntactic, spatial, and semantic features. First, a *table* is a two-dimensional arrangement of $n \times k$ items, so that the position of each cell of the table can be uniquely indexed by a pair $\langle x, y \rangle$ of positive integers, with $x \leq n$ and $y \leq k$. Second, tables are presented spatially as horizontal *rows* and vertical *columns*. Third, it is a characteristic feature of a table that its rows and columns exhibit some *semantic unity*, i. e., that each row and column can be understood as forming a meaningful entity.[1] An arrangement of cells into rows and columns that lacks this kind of semantic unity is an *array* or a *grid* [8]. Examples of particular representations that make use of columns (or rows), but lack the semantic unity of their rows (or columns, respectively), or that are simply grids are abaci, bar charts, and cellular automata.

Syntactic and Spatial Features of Tables. According to the definition given above, a table is an array of cells, such that some of them are adjacent (*neighbors*); the individual indices of neighboring cells differ only by 1. Figure 1 contains an example of a table showing the indices of each cell (on the left) and the relative indices of the neighbors of the cell $\langle x, y \rangle$.[2] Due to the spatial structure of a table, the cells are typically presented in a rectangle and we can easily determine the rows and columns, as well as various diagonals: cells, whose indices that have the same x-coordinate form a row, those with the same y-coordinate a column.

[1] Such collections of multiple components of a diagram that 'form a unit with semantic significance' are called 'global objects' in [26, p. 261].

[2] The corner of the table in which cell $\langle 1, 1 \rangle$ is positioned is arbitrary.

$$\begin{array}{cccc} \langle 1,1\rangle & \langle 1,2\rangle & \ldots & \langle 1,k\rangle \\ \langle 2,1\rangle & \langle 2,2\rangle & \ldots & \langle 2,k\rangle \\ \vdots & \vdots & \ddots & \vdots \\ \langle n,1\rangle & \langle n,2\rangle & \ldots & \langle n,k\rangle \end{array} \qquad \begin{array}{ccccc} \langle x-1,y-1\rangle & & \langle x-1,y\rangle & & \langle x-1,y+1\rangle \\ & \nwarrow & \uparrow & \nearrow & \\ \langle x,y-1\rangle & \leftarrow & \langle x,y\rangle & \rightarrow & \langle x,y+1\rangle \\ & \swarrow & \downarrow & \searrow & \\ \langle x+1,y-1\rangle & & \langle x+1,y\rangle & & \langle x+1,y+1\rangle \end{array}$$

Fig. 1. Indices of an $n \times k$ table (on the left) and indices of neighboring cells.

Two cells that differ by 1 in both of their coordinates are diagonal. It is a crucial aspect of tables that their syntactic features are directly related to their spatial features, which will be relevant in our discussion in Sect. 2.

In many cases, the presentation of a table is augmented by *labels* for the rows and columns, which can be used either to denote the rows and columns, or to convey information about the type of their content. Particular presentations of tables frequently also make use of horizontal or vertical lines, coloring or shading of alternate rows or columns, and different spacing between rows and columns to emphasize the semantic unity of the entries and to guide the reading of the table [28, p. 81–95].

Semantic Features of Tables. We can also distinguish certain types of tables on the basis of the content of their cells: in a *homogeneous* table all cells are of the same type (e. g., numbers); in a *relational* table all cells can take only two values, such as empty/nonempty, 0/1, or true/false. Such tables are also called 'feature tables' and are discussed in [19]. More important for their practical use, however, is the *ordering* of the data in some of the rows or columns. Accordingly, Wainer's first rule for the preparation of useful tables is: 'Order the rows and columns in a way that makes sense' [31]. Indeed, hardly any tables we come across do not exhibit some kind of ordering of their rows or columns: ledgers are ordered by date, timetables by time, directories by last name, duty rosters by time and name, inventory lists by item codes, etc. This is a semantic feature, because it does not depend on the spatial arrangement of the cells, but on their content. The ordering (or *sorting*) of a table implies a certain direction of reading ('directedness' is one of the basic aspects of the 'grammar' of diagrams in [11, p. 71–73]), which in turn has the effect that a particular cell is identified as the *origin*, which is typically where one would start reading the entries. As the example of a Japanese train schedule reproduced in [28, p. 47] shows, the origin does not have to be in one of the corners of the table. The rows and columns, together with the origin and directions of reading constitute what Tufte calls the 'viewing architecture' of a table [29, p. 159].

Effects on Perception and Cognition. Since tables are often perceived as a whole, Gestalt psychology provides us with a useful framework for their study. In particular, we will appeal to the Gestalt principles of similarity, proximity, good continuation, and symmetry [5, 17]. The spatial arrangement of a table immediately makes us perceive the rows and columns as individual entities. For example, if the columns have entries of different types, e. g., text, dates, numbers, then we tend to perceive these entries as belonging together (i. e., as forming a column), due to the Gestalt principle of *similarity*. By the same principle, coloring can be used so that we perceive the items as belonging together and thereby forming a unit. In addition, we can also employ judicious spacing: if the rows are closer together than the columns, then, according to the Gestalt principle of *proximity*, we perceive the items above each other as being connected, i. e., the columns stand out, and vice versa. This phenomenon is illustrated by the following arrangements:

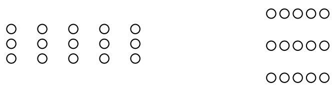

Not only does the presentation of cell items guide our perception towards the rows and columns, but it also suggests certain axes for reading and moving through the table. That the ordering of entries suggests a particular starting cell (origin) and a direction for reading was already mentioned above. The proximity of neighboring cells similarly suggests a movement in horizontal, vertical, and diagonal directions (indicated by the arrows in Fig. 1, right), and the Gestalt principle of *good continuation* or *continuity of direction*, then, continues to guide our perception along that line.

Materiality. Timetables, which are often adduced as prototypical examples of tables are static, printed on paper or displayed on a computer screen. This limits our possible interactions with them to a mainly passive use of retrieving information. Nevertheless, the particular form and kind of material makes a difference even in this case, e. g., when the table extends over many pages. Moreover, we can also produce tables ourselves, e. g., ledgers, which allows us to further extend and manipulate them. Spreadsheet software allows even further possibilities, such as including formulas that refer to the values of other cells and sorting the entire table according to the entries of a specific row or column.

3 Access to Global Information

After the introduction of the main features of tables, we shall now address how these features contribute to the use of tables as cognitive tools. We begin with the access to global information, before turning to different algorithms for accessing particular cells (Sect. 4) and the visualization of relational patterns (Sect. 5).

Early examples of the arrangement of data into tables are *ledgers*, e. g., to record the sales in a store [8, p. 125–146].[3] Historically, such ledgers often consisted of sheets that could be spread on a table (their modern cousins, called 'spreadsheets', are discussed below, in Sect. 6). In a typical ledger, each row represents an individual sale, recording the date, item, buyer, and price. In addition to simply recording the information, the tabular arrangement offers some very specific ways of assessing it:

[3] See [1] for many examples of the use of tables in the history of mathematics.

(1) It affords a quick *overview* of the data,

e. g., regarding the total number of sales, the distribution of customers and prices. Such an overview also allows for

(2) the identification of *singularities*,

such as particularly large numbers in a place-value system, or empty cells in an otherwise filled table; they stand out, because they violate the Gestalt principle of *similarity* and are therefore immediately perceived. Thus, in practical terms,

(3) such an arrangement makes it easier to *detect missing data*

(as the corresponding cells are empty) than a linear presentation. Moreover,

(4) it provides for means to check for the *correctness* of the data,

e. g., dates must have a fixed form; cells in a number-column cannot contain letters; numerical values within a column are typically allowed only within a certain range; if a column is sorted, violations of the order can be detected. Note, that these checks can be performed purely syntactically, i. e., without knowledge of the specific meanings of the entries.

4 Information Retrieval

In discussions about the advantages of tabular representations, it is frequently claimed, on the basis of an appeal to the reader's intuition, that they make it easier to locate information (see, e. g., [6, p. 80]). Consider the example of a timetable, such as the following departure times at a bus station:

8:00	8:20		8:40
9:00		9:30	
10:00	10:20		10:40

Clearly, this information could be represented linearly, but with a tabular representation it is easy to find the next departure time by first finding the appropriate row on the basis of the entries in the left-most column and then going through the columns. Thus, in addition to the quick overview of the data that is provided by a table, we have identified another important advantage of tables, namely the

(5) quick *access* to specific cells.

This property underlies the long-lasting use of tables to represent astronomical data in science and logarithmic tables, as well as addition and multiplication tables in mathematics.

Before moving on, we should pause for a moment and reflect once more on the previous claim. After all, *binary search* is in general the best algorithm for searching for an entry in an ordered list, so let us briefly look at how the representation of the data interacts with the complexity of search algorithms and whether tables really offer an advantage over linear representations.

To find an element in a sorted list consisting of $n \cdot k$ elements using binary search takes $\log_2(n \cdot k)$ steps in the worst case. Now, if these elements are arranged in a $n \times k$ table, applying a binary search to the rows takes $\log_2 n$ steps and applying it again to the columns in that row takes another $\log_2 k$ steps, resulting in a total of $\log_2 n + \log_2 k$ steps. This, however, is equivalent to $\log_2(n \cdot k)$. Thus, when using binary search, it makes no difference whether the data is presented as a table or linearly, so we cannot explain the intuition behind the efficacy of tables for information retrieval on the basis of this search algorithm. (We analyzed only the worst case scenario, but considerations about average cases are analogous.)

While considerations of computational complexity provide a useful means of comparison between different representations, if we want to use them to tell us something about the efficacy of representations for *human* use, we must also take into account the plausibility of the assumptions on which they rely, in this case that each computational step has the same cost. In practice, however, this assumption is frequently unjustified. Consider, for example, the task of searching for a particular name in an old-fashioned telephone book. Given that the entries are sorted by name, we could do a binary search: Open the book in the middle and determine whether the name we are looking for occurs in the first half or in the second; then, go to the relevant half and repeat the procedure, by looking again at the page in the middle of that half of the book, and so on. Instead of applying this procedure, however, what many people will do is to flip through the book until one reaches pages whose entries begin with the same letter as the name we are looking for. Then, one flips back or forth page by page, depending on where one expects the name to be. In terms of complexity theory, this procedure corresponds more or less to a *linear search* through the phone book, which in general involves vastly more steps than a binary search. However, the physical operation of 'flipping through the pages' might well be easier and quicker than 'open the book in the middle of a given block of pages', which would be a good reason for many people's preference for using a linear search algorithm.

So, if we use a linear search algorithm, how does the presentation of the data affect the complexity of the search? A linear search in a sorted list of $n \cdot k$ elements takes $n \cdot k$ steps in the worst case. However, a linear search of the rows takes at most n steps and a linear search of the columns another k, resulting in a total of $n + k$ steps. Here, the difference is considerable, as $n + k$ is usually much smaller than $n \cdot k$. (In a $n \times n$ table, the difference is between $2n$ and n^2 steps, i. e., between linear and quadratic complexity classes.) In other words, when performing a linear search (which presumably comes more naturally to most of us), using a table to represent the data yields much faster search results, just as our intuitions predicted.

In addition to the material conditions that can affect the ease with which certain operations can be carried out, a linear search algorithm might also be preferred because it exploits features of our perceptual apparatus, which allows the scanning of individual elements of a list at a glance (in particular if the list is small), whereas determining the middle element of a list would require additional reasoning. In fact, for lists of relatively short length the actual differences in necessary steps are fairly small and complexity considerations become relevant only with large numbers of elements. Moreover, to rows or columns that are not sorted, as is the case for ledgers and many other tables, binary search cannot be applied at all. Thus, with regard to locating elements, there is a considerable interaction between physical, perceptual, and cognitive constraints, with the content and structural features of the representation.

5 Visualization of Relational Structure and Patterns

As I have argued above, the spatial arrangement of a table allows us to easily perceive some of the items as belonging together (in particular, those arranged in rows, columns, and diagonals). Furthermore, *symmetries* stand out frequently, too, as the recognition of symmetries is also a Gestalt principle. In short, tables

(6) facilitate the perception of particular *patterns* in the data.

If these patterns also correspond to meaningful properties of the represented subject matter, this amounts to an immediate visualization of structural information contained in the data.

Even in the absence of deep structural relations, merely displaying data in a table can focus a researcher's attention to a connection between the ordering of the table and particular features of the data. For example, such connections were found in the tables used for the classification of chemical substances by Doumas and Boullay (1828), which is discussed by Klein in order to illustrate the role of 'paper tools' in the creation of scientific knowledge [9]. Polya quotes Jacob Bernoulli (1713) as noting that 'This table of numbers has eminent and admirable properties'; in his book on problem solving, Polya himself frequently suggests to arrange mathematical formulas 'in suitable tables' [16, vol. 2, p. 193] and encourages the student to ask: 'Do you notice something worth noticing—some law or pattern or regularity?' [16, vol. 2, p. 152].[4]

In order to account for the fact that images and graphs can represent information at various levels of abstraction in a way that makes it immediately available to the users, Kulvicki introduced a notion of 'extractability' of information that is based on a correlation between what is 'syntactically' and what is 'semantically salient' [12]. As we shall see in the following example of tabular representations of binary operations, the specific (syntactic) patterns are straightforwardly recognizable and their correlation to (semantic) properties of the operations that are represented can be easily established. Thus, Kulvicki's analysis of images (as opposed to linguistic representations) seems to apply to tables just as well; indeed, in most examples that he gives himself, the information is actually represented in a table.

To illustrate the previous claim, let us look at some simple examples in which tables that represent the results of binary operations contain easily identifiable

[4] Of course, patterns can also emerge from arrangements that are not tables, as Ulam's famous patterns of the distribution of prime numbers show [22].

$+_3$	0	1	2
0	0	1	2
1	1	2	0
2	2	0	1

$\times_3$	0	1	2
0	0	0	0
1	0	1	2
2	0	2	1

$\rightarrow$	T	F
T	T	F
F	T	T

Fig. 2. Tables for addition and multiplication modulo 3, truth table for implication.

syntactic patterns that correspond to properties of the operations: the results of addition and multiplication modulo 3 are shown in the first two labeled homogeneous tables in Fig. 2. Here some observations about the distribution of the values in the tables: (i) Each entry consists of one of the values 0, 1, 2, which are also the input values for the operations according to the labels; this corresponds to the fact that the operations are *closed*, i. e., they do not return any value that is not among the inputs. (ii) The fact that the top line of the $+_3$-table corresponds exactly to the labels indicates that the element 0 is the *left-identity* of the operation, just as 1 is the *left-identity* of $\times_3$. These elements are also right-identities, as the elements in their columns match exactly the column labels. These observations together establish 0 as the identity element for $+_3$, and 1 as the identity element for $\times_3$. (iii) Because the identity element for $+_3$ occurs in each row and each column, the operation has an *inverse*, which is not the case for $\times_3$. These observations identify $+_3$ as a *group* operation. (iv) The fact that the entries are symmetrical with respect to the main diagonal from top left to bottom right indicates that both operations $+_3$ and $\times_3$ are *commutative*. Applying the same analysis to the truth-table for classical implication (shown in Fig. 2, right) yields that the operation is also closed, has a left-identity (T), but no right-identity, has no inverse operation, and is not commutative.

Tables that exhibit particular patterns are also easier to memorize, as the historian of mathematics Swetz remarks: 'Multiplication facts were organized into tables, traditionally called *Pythagorean tables*, in which the numerical patterns might be better observed and remembered' [25, p. 84].

6 Operations on Tables

We turn now to the ways in which one can actively engage with tables by performing operations on cells or groups of cells.

Functions on Cells. The notion of 'derivative meaning' was introduced by Shimojima as 'the additional informational relation derivable from, but not included in, the system's basic semantic conventions' of a representation [20, p. 114]. For the case of tables, the basic semantic conventions are those that give the meanings of the values in individual cells. Typical derived information that can be obtained in a table involves those values that are contained in entire rows and columns, for example: how many elements are contained in a row or column, the sum of the elements, the average, or other statistical or arithmetical functions.

Name	T1	T2	T3
A	50	75	85
B	30	50	55
C	40	85	65

Name	T1	T2	T3	Sum
A	50	75	85	210
B	30	50	55	144
C	40	85	65	180
Avg.	40	70	68	

Fig. 3. Simple tabular arrangement of test results with derived meanings.

As a simple illustration, consider the results of three tests (T1, T2, and T3) of three students (A, B, C), as shown in the left table in Fig. 3. This table is used to derive additional information, namely the sums of the results for each student and the averages of each test, recorded in the table on the right. We can readily see that the average of the first test is considerably lower than the other two, and that student A had the highest overall score. I take it that this use of tables is familiar and hardly surprising, but the question arises: *why* do tables lend themselves to this usage? The answer brings us back to the features of tables introduced in Sect. 2. Due to the semantic unity of the rows and columns, to determine the average result of a test or sum of the results of an individual student, only the values in a single column (or row, respectively) have to be consulted; the spatial arrangement makes this task easier, because all cells that have to be taken into consideration are direct vertical (or horizontal, respectively) neighboring cells. This simplifies the symbolic expression that is used to refer to the range of these cells (when formulas on the cells are entered in a spreadsheet) and facilitates shifting the focus of one's attention (in case the derived meanings are obtained 'by hand'). Thus, it is the combination of the syntactic, spatial, and semantic characteristics of tables, together with the effect these have on our perception, that underlies the ease in which derived information can be obtained from a tabular representation.[5]

Spreadsheets. Operations on tables that can be expressed as functions on cells are an essential feature of spreadsheets. However, in addition to simply automating operations that could otherwise also be done by hand, spreadsheets also offer the ability of sorting a table according to the values in a row or column. To achieve this with traditional tables, one would have to create a new, separate table. The fact that spreadsheets allow for in-place sorting, or, to use Stenning's term, an 'agglomerative' mode of reasoning [23, p. 41], puts some pressure on considering the difference between tables and spreadsheets as merely one pertaining to their materiality. Rather, it seems that tables and spreadsheets are

[5] Turing's general analysis of computation also begins with a two-dimensional writing surface: 'Computing is normally done by writing certain symbols on paper. We may suppose this paper is divided into squares like a child's arithmetic book. In elementary arithmetic the two-dimensional character of the paper is sometimes used' [30, p. 239]. Although, Turing does not consider this to be an essential feature of computations and thus restricts his model to a one-dimensional tape. Nevertheless, a Turing Machine can also move its head only to adjacent cells on the tape.

really two different kinds of representations. However, we must leave this issue as a topic for future discussion.

Matrices. The study of equations in terms of tabular arrangements of their coefficients was practiced already in ancient China [2], but matrices, as genuine mathematical objects became popular only in the nineteenth century. While they can be understood generally as arrays of numbers, in many applications their rows and columns represent clearly discernible semantic units. Moreover, although many operations on matrices, such as transposition and matrix multiplication, are formally defined in terms of the indices of the elements, they are often conceived, taught, and memorized in terms of operations on rows and columns, thereby appealing directly to the spatial structure of matrices.

7 Tabular Manipulations of Structured Notations

In the previous sections I have argued that the overlapping of the syntactic and spatial features of tables with the semantic unity of the rows and columns allows for, among other things, an immediate perception of patterns and a straightforward formulation of operations on the data. So far, however, we have considered cells as containing only basic elements, e. g., linguistic and numeric items. We come now to a new aspect of the power of tables as representational tools, namely in the case that the cells contain only parts of more complex expressions.

Due to their recursive grammar and compositional semantics, expressions in most mathematical notations are themselves structured. Because of this, we can not only use tables to represent the relations between different expressions, but also to represent the relations of the symbols that constitute an expression. Take numerals, for example. We have already seen examples where numerals were elements in a table, but we can also use a tabular arrangement for the numerals themselves. For the sake of illustration, let us first consider a more exotic example, namely Roman numerals, and then turn to more familiar algorithms for the decimal place-value system and algebraic equations.

Roman Numerals. In the left table of Fig. 4 we see the standard, linear representation of two Roman numerals in the purely additive format. We can perceive at a glance that the first one takes up more space, thus presumably has more symbols, but because of the different width of the symbols ('M' vs. 'I'), we cannot be completely sure about that. On the right side of the figure the same numerals are represented in a tabular format, in which each row contains a numeral expression and there is a column for each group of letters. Here we immediately

Linear representation	Tabular representation						
MDCCCLXXII	M	D	CCC	L	XX		II
MMCXXXVI	MM		C		XXX	V	I

Fig. 4. Linear and tabular representations of Roman numerals.

recognize that the second numeral has two M's and thus represents a greater value than the first. The readability of the numerals is greatly enhanced by the tabular representation. Since additions of such numerals are done letter-by-letter (with subsequent simplifications), adding numerals in this format can be done column-wise, thus simplifying the procedure considerably. Indeed, in their discussion of addition and multiplication algorithms for Roman numerals, Schlimm and Neth use such a tabular representation [18].

Addition and Multiplication Algorithms. Most likely we are so familiar with our paper-and-pencil algorithms using the Indo-Arabic decimal place-value system that we tend to overlook the basic assumptions that make them possible in the first place. In particular, most addition algorithms presuppose that the numerals are written underneath each other in a right-aligned way (compare the additions shown in Fig. 5). In fact, if the numerals were left-aligned or not written underneath each other at all, the formulation of an addition algorithm would be much more complex, since it cannot rely on the column-wise processing of the digits. Thus, the (correct) tabular arrangement of the numerals is crucial for our familiar, simple, and efficient addition algorithm.

The tabular arrangement of the addends also plays a crucial role in common multiplication algorithms. Notice the placement of the intermediate results of the addends shown in the first two multiplications in Fig. 6. Here, '115' and '92' are not right-aligned, because they are obtained by multiplication of single digits that have different power-10 factors. In this way, however, the usual column-wise addition algorithm can still be applied. (Historically, this algorithm was known as 'chessboard multiplication', clearly in reference to the tabular arrangement [24, p. 205].) In the third multiplication shown in Fig. 6, the intermediate results are right-aligned, which has the consequence that we have to add the digits *diagonally*, instead of column-wise (which is indicated by the shading of the diagonal). An even more refined algorithm, sometimes called 'lattice multiplication', can be found in the *Treviso Arithmetic* of 1478 [24] (shown in the fourth example in Fig. 6). Here, no intermediate handling of carries is necessary, as the intermediate results of the single-digit multiplications are simply written out, which requires more diagonals to be considered when adding these intermediate results. What is important for our discussion is that all of these algorithms make *essential use* of the fact that a tabular arrangement of the intermediate results allows for a straightforward computation of the final result by processing single digit addition either column-wise or diagonally.

$$\begin{array}{r} 9\;2\;1 \\ +1\;1\;5\;3 \\ \hline 2\;0\;7\;4 \end{array} \qquad \begin{array}{l} \;\;\;\,9\;2\;1 \\ +1\;1\;5\;3 \\ \hline \;\;2\;0\;7\;4 \end{array} \qquad 921 + 1153 = 2074$$

Fig. 5. Simple additions with Indo-Arabic numerals.

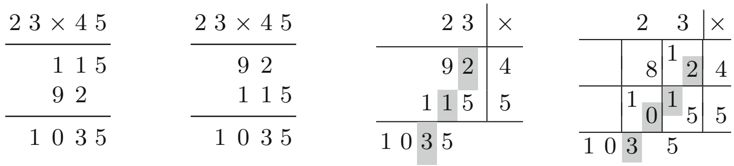

Fig. 6. Simple multiplications with Indo-Arabic numerals.

Algebraic Equations. To counter the impression that the advantages of the use of tables for structured notations applies only to numerals, let us briefly look at another example, namely algebraic equations. Consider the equations $x^2 + 3x = 2$ and $y^2 = -3y + 2$. A reader with some mathematical experience might be able to parse them quickly, but, in general, the relations between the two equations are easier to detect if they are presented as follows:

$$\begin{array}{llcll} x^2 & +3x & = & & 2 \\ y^2 & & = & -3y & +2 \end{array}$$

The tabular form of this presentation is determined by the following features: Each of the two rows represents an equation and even without any labels we can readily discern the columns, organized in terms of the powers of the variables, the signs for arithmetical operations, and the equality symbol. By scanning the columns we quickly realize the following differences: In the first column, the elements differ only in the names of the variables ('x' and 'y'), but not in their power; a term corresponding to '$3x$' is missing in the second equation, while a term corresponding to '$-3y$' is missing in the first equation; these two terms occur on different sides of the column with the equation sign and they differ in their leading sign ('+' vs. '−'). With the additional knowledge that a term can be 'pushed' to the other side of the equality symbol while reversing the leading sign, we can *see* that the two equations express the same condition for the free variable and that they only differ in the particular name of this variable ('x' and 'y'). These considerations about the manipulation of algebraic equations are in accord with the 'perceptual account of symbolic reasoning' [13], which is mainly based on empirical work by Landy and Goldstone. In fact, examples given to illustrate this account use a the tabular representation of equations (see, e. g., Fig. 1 in [14, p. 1073]).

8 Operations on Infinite Tables

So far we have considered only finite, two-dimensional tables, but some of their characteristic features remain in place if we generalize the concept to infinitely many rows and/or columns, as will be illustrated in the following two proofs.[6]

[6] Extending tables to three or more dimensions is beyond the scope of this paper.

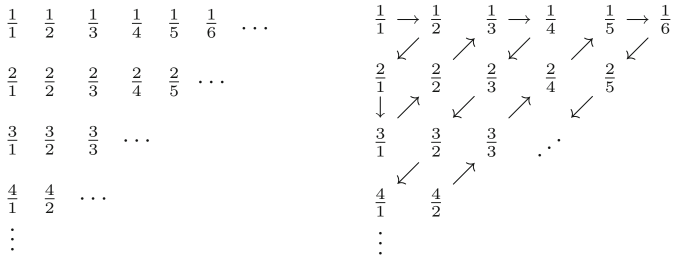

Fig. 7. Tables illustrating the argument for the denumerability of the rationals.

Dovetailing. It is a surprising fact that follows from Cantor's definition of infinite cardinalities, which asserts that two sets have the same cardinality if and only if they can be put into a 1-1 correspondence, that the set of rational numbers has the same cardinality as the set of natural numbers. To establish this, one has to produce a 1-1 correspondence between them, such as

1	2	3	4	5	6	...
$\frac{1}{1}$	$\frac{1}{2}$	$\frac{2}{1}$	$\frac{3}{1}$	$\frac{2}{2}$	$\frac{1}{3}$	...

An untrained reader might wonder what exactly the rule is that establishes the correspondence (i. e., which rational numbers follow in the sequence and are correlated to 7, 8, and 9?) and whether this table actually contains *all* fractions. A different arrangement of the fractions gives an answer to these questions in an intuitive way. The table on the left side of Fig. 7 shows an arrangement of the rational numbers, represented as fractions, and the table on the right illustrates how one can count the fractions, i. e., establish a 1-1 correspondence between the natural numbers and the rational numbers. (In fact, the table also shows that some rational numbers are counted more than once, since the fractions $\frac{1}{1}$ and $\frac{2}{2}$, for example, represent the same rational number, so additional restrictions must be imposed to make this work).

Diagonalization. The well-known presentation of the argument that the real numbers between 0 and 1 are *uncountable*, i. e., that it is impossible to establish a 1-1 correlation between them, begins by assuming that it is possible to have a 1-1 mapping f between the natural numbers and the real numbers in question. This is illustrated by a table in which we have the natural numbers as labels for the rows and each element of the decimal expansion of the corresponding real numbers in one column each (as shown in Fig. 8, left; this is another example of a tabular representation of a structured notation, discussed in the previous section). The construction of a real number that cannot be in this table is obtained by changing the first element in the decimal expansion of the first real number in the table, the second element of the expansion of the second number, and so on; these elements are marked in the table on the right of Fig. 8 by boxes. The name of this proof technique, 'diagonalization', clearly refers to the tabular arrangement shown and the fact that this presentation has

ℕ	ℝ						
1	0.	1	2	3	4	5	...
2	0.	5	0	0	0	0	...
3	0.	3	3	3	3	3	...
4	0.	1	2	1	2	1	...
⋮	⋮						⋱

ℕ	ℝ						
1	0.	$\boxed{1}$	2	3	4	5	...
2	0.	5	$\boxed{0}$	0	0	0	...
3	0.	3	3	$\boxed{3}$	3	3	...
4	0.	1	2	1	$\boxed{2}$	1	...
⋮	⋮						⋱

Fig. 8. Part of the argument of the uncountability of ℝ.

become the default in textbooks on set theory can be taken as evidence for its pedagogical value (e. g., [4, p. 132]; for the historical development of Cantor's proofs, see [27]).

9 Discussion

Tables as Effective Representations. Because of the semantic unity of the rows and columns of a table, any element is (semantically) related to other elements in two different ways, which also correspond to the syntactic (rows and columns) and spatial arrangement of the elements. Thus, when we see two elements next to each other, we also 'see' their semantic connection: syntactic movement (e. g., in a row) corresponds to spatial movement (e. g., to the right) and the realization that part of the meaning of the element stays the same and part of it changes. This alignment of syntax, semantics, and visual appearance is what makes tables so versatile and powerful representations. For example, we have identified the following aspects of tables in our discussion: (1) They afford a quick overview of the data, (2) allow for the identification of singularities, (3) make it easier to detect missing data, (4) provide for means to check for the correctness of the data, (5) provide quick access to specific cells, and (6) facilitate the perception of particular patterns in the data. Moreover, tables facilitate the expression of functions on their elements and can be used to simplify the manipulations of structured notations.

What is involved in retrieving the information represented in a table is deeply embedded in our reading habits along vertical and horizontal axes (from left to right, from top to bottom) and our background knowledge of ordered sequences (alphabet, numerals, cultural conventions such as the order of first name/last name). The Gestalt principles of proximity, similarity, good continuation, and symmetry, underlie our spontaneous perceptual and cognitive reactions to tabular representations. Due to the two-dimensional organization, binary relations that hold between the elements appear as visual patterns that can be immediately recognized. In this way, tables can be said to translate abstract relations into a perceivable form. If structural relations hold between the components of expressions in some notation, then a tabular arrangement can support the understanding of these expressions as well as their manipulations. Just as formulas can be productive tools in science [9] and numerals can be tools for operating with numbers [10], tables are a powerful tool for our processing of information.

Tables and Diagrams. One motivation to take a closer look at tables, to determine their characteristic features, and to study how these contribute to the efficiency of this kind of representation, has been their relative simplicity and the fact that they are rather constrained. This lack of generality was offset by a clear identification of their structural (syntactic and spatial) and semantic features, and how these affect our perception and our ability to reason with them. If *diagrams* are understood from a semiotic point of view as structured signs, then tables are simply a particular, well-defined category of diagrams. However, diagrams that are frequently discussed in the literature (e. g., [3,6,15,20]) are often *more general* than tables (pace Stenning, who considers directly interpreted diagrams to be subject to more constraints than tables [23, p. 45]), since they do not require any vertical or horizontal arrangement of their elements, and as *more expressive*, since they typically use lines or arrows to represent relations between elements. Nevertheless, it also seems that, whenever possible, such diagrams are presented in such a way that they look very much like tables! The most striking examples are commutative diagrams, in which the elements are arranged in rows and columns in such a way that these also exhibit a semantic unity (see, e. g., [20, p. 122], [3, p. 3,18]). This suggests that it might be fruitful to use the characteristic features of tables identified above and their effects on perception and cognition also as ingredients of a more general discussion of diagrams.

Acknowledgments. Many thanks to Alberto Lopez Garcia, Tyler Marghetis, Valérie Therrien, Theresa Wege, and three anonymous reviewers for comments on an earlier draft. This research was supported by the Social Sciences and Humanities Research Council of Canada.

References

1. Campbell-Kelly, M., Croarken, M., Flood, R., Robson, E. (eds.): The History of Mathematical Tables. From Sumer to Spreadsheets. Oxford University Press, Oxford (2003)
2. Chemla, K.: Different concepts of equations in The Nine Chapters on Mathematical Procedures and in the commentary on it by Liu Hui (3rd century). Hist. Sci. **4**(2), 113–137 (1994)
3. De Toffoli, S.: What are mathematical diagrams? Synthese **200**(2), 1–29 (2022)
4. Enderton, H.B.: Elements of Set Theory. Academic Press, New York (1977)
5. Enns, J.T.: Perception, Gestalt principles of. In: Nadel, L. (ed.) Encyclopedia of Cognitive Science. Nature Publishing, London (2003)
6. Giardino, V.: Behind the diagrams: cognitive issues and open problems. In: Krämer, S., Ljungberg, C. (eds.) Thinking with Diagrams. The Semiotic Basis of Human Condition, pp. 77–101. De Gruyter Mouton, Boston (2016)
7. Guthrie, J.T., Weber, S., Kimmerly, N.: Searching documents: cognitive processes and deficits in understanding graphs, tables, and illustrations. Contemp. Educ. Psychol. **18**, 186–221 (1993)
8. Higgins, H.B.: The Grid Book. The MIT Press, Cambridge (2009)

9. Klein, U.: Paper tools in experimental cultures. Stud. Hist. Philos. Sci. **32**(2), 265–302 (2001)
10. Krämer, S.: Writing, notational iconicity, calculus: On writing as a cultural technique. Modern Language Notes - German Issue **118**(3), 518–537 (2003)
11. Krämer, S.: Figuration, Anschauung, Erkenntnis. Grundlinien einer Diagrammatologie. Suhrkamp, Berlin (2016)
12. Kulvicki, J.: Knowing with images: medium and message. Philos. Sci. **77**(2), 295–313 (2010)
13. Landy, D., Allen, C., Zednik, C.: A perceptual account of symbolic reasoning. Front. Psychol. **5**(275), 1–10 (2014)
14. Landy, D., Goldstone, R.L.: How much of symbolic manipulation is just symbol pushing? In: Proceedings of the 30st Annual Conference of the Cognitive Science Society, Austin, TX, pp. 1072–1077. Cognitive Science Society (2009)
15. Larkin, J.H., Simon, H.A.: Why a diagram is (sometimes) worth ten thousand words. Cogn. Sci. **11**, 65–99 (1987)
16. Polya, G.: Mathematical Discovery. On Understanding, Learning, and Teaching Problem Solving, Combined edn.. Wiley (1981). first ed. 1962–65
17. Rock, I., Palmer, S.: The legacy of gestalt psychology. Sci. Am. **263**(6), 84–90 (1990)
18. Schlimm, D., Neth, H.: Modeling ancient and modern arithmetic practices: Addition and multiplication with Arabic and Roman numerals. In: Sloutsky, V., Love, B., McRae, K. (eds.) Proceedings of the 30th Annual Meeting of the Cognitive Science Society, Austin, TX, pp. 2097–2102. Cognitive Science Society (2008)
19. Shimojima, A.: The inferential-expressive trade-off: a case study of tabular representations. In: Hegarty, M., Meyer, B., Narayanan, N.H. (eds.) Diagrams 2002. LNCS (LNAI), vol. 2317, pp. 116–130. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-46037-3_16
20. Shimojima, A.: Semantic Properties of Diagrams and Their Cognitive Potentials. CSLI Publications, Stanford (2015)
21. Shin, S.J.: The Logical Status of Diagrams. Cambridge University Press, Cambridge (1995)
22. Stein, M.L., Ulam, S.M., Wells, M.B.: A visual display of some properties of the distribution of primes. Am. Math. Mon. **71**(5), 516–520 (1964)
23. Stenning, K.: Seeing Reason: Image and Language in Learning to Think. Oxford University Press, Oxford (2002)
24. Swetz, F.: Capitalism and Arithmetic: The new math of the 15th century, including the full text of the Treviso arithmetic of 1478. Open Court, La Salle, Ill (1987)
25. Swetz, F.: To know and to teach: mathematical pedagogy from a historical context. Educ. Stud. Math. **29**(1), 73–88 (1995)
26. Takemura, R., Shimojima, A., Katagiri, Y.: Logical investigation of reasoning with tables. In: Dwyer, T., Purchase, H., Delaney, A. (eds.) Diagrams 2014. LNCS (LNAI), vol. 8578, pp. 261–276. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-44043-8_27
27. Therrien, V.: The evolution of Cantor's proofs of the non-denumerability of $\mathbb{R}$. Manuscript (2022)
28. Tufte, E.R.: Envisioning Information. Graphics Press, Cheshire (1990)
29. Tufte, E.R.: The Visual Display of Quantitative Information, 2nd edn. Graphics Press, Cheshire (2001)
30. Turing, A.M.: On computable numbers, with an application to the Entscheidungsproblem. Proc. Lond. Math. Soc. **42**, 230–265 (1936)
31. Wainer, H.: Understanding graphs and tables. Educ. Res. **21**(1), 14–21 (1992)