REPRESENTATIONS in MIND and WORLD

Essays Inspired by Barbara Tversky

Edited by Jeffrey M. Zacks and Holly A. Taylor

ROUTLEDGE

REPRESENTATIONS IN MIND AND WORLD

Essays Inspired by Barbara Tversky

"This remarkable volume celebrates and documents the accomplishments of a remarkable cognitive scientist, Barbara Tversky. In their investigations of psychological representation Barbara and her students elegantly bridge at least three formidable gulches: between mental and physical representation, between lab and field, and between disciplines that not only include philosophy and psychology but also the visual arts, design, and computer science.

This work not only honors George Miller's exhortation to give psychology away, but closes the circle by using an analysis of physical representation in the field to shape and inform psychological theory. Each chapter shows the benefits and wisdom of establishing a permanent dialogue between academic theory and phenomena and challenges in the real world, creating not 'applied' science but science that is richer and more complete by virtue of worrying about its applicability."

—**Elke U. Weber**, Princeton University

"Barbara Tversky is our premier cognitive ethologist. She observes cognition at work in the wild, then brings it into the lab. This illuminating Festschrift shows the range and power of the work now being done by her numerous students and collaborators. I recommend it strongly."

—**C. R. Gallistel**, Rutgers University—New Brunswick

This volume pulls together interdisciplinary research on cognitive representations in the mind and in the world. The chapters—from cutting-edge researchers in psychology, philosophy, computer science, and the arts—explore how structured representations determine cognition in memory, spatial cognition information visualization, event comprehension, and gesture. It will appeal to graduate-level cognitive scientists, technologists, philosophers, linguists, and educators.

Jeffrey M. Zacks is Professor and Associate Chair of Psychological & Brain Sciences, and Professor of Radiology, at Washington University in St. Louis.

Holly A. Taylor is Professor of Psychology and Co-Director of the Center for Applied Brain & Cognitive Sciences at Tufts University.

REPRESENTATIONS IN MIND AND WORLD

Essays Inspired by Barbara Tversky

Edited by Jeffrey M. Zacks and Holly A. Taylor

Routledge
Taylor & Francis Group
NEW YORK AND LONDON

First published 2018
by Routledge
711 Third Avenue, New York, NY 10017

and by Routledge
2 Park Square, Milton Park, Abingdon, Oxon, OX14 4RN

Routledge is an imprint of the Taylor & Francis Group, an informa business

Library of Congress Cataloging-in-Publication Data
A catalog record for this book has been requested

ISBN: 978-1-138-82969-5 (hbk)
ISBN: 978-1-138-82970-1 (pbk)
ISBN: 978-1-315-16978-1 (ebk)

Typeset in Bembo
by Keystroke, Neville Lodge, Tettenhall, Wolverhampton

Paperback cover image courtesy Andrea Kantrowitz

CONTENTS

LIST OF CONTRIBUTORS

Kathleen M. Arnold, Department of Psychology & Neuroscience, Duke University, United States.

Thomas Barkowsky, Bremen Spatial Cognition Center, University of Bremen, Germany.

Mireille Bétrancourt, School of Education & Technology, University of Geneva, Switzerland.

Tad T. Brunyé, Center for Applied Brain & Cognitive Sciences at Tufts University; Department of Psychology, Tufts University; and Cognitive Science Team, U.S. Army Natick Soldier Research, Development, and Engineering Center, United States.

Roberto Casati, Institut Jean Nicod CNRS ENS EHESS, France.

Michel Denis, LIMSI-CNRS, Université Paris Sud, France.

Frank Dylla, Bremen Spatial Cognition Center, University of Bremen, Germany.

Zoe Falomir, Bremen Spatial Cognition Center, University of Bremen, Germany.

Nancy Franklin, Department of Psychology, Stony Brook University, United States.

Christian Freksa, Bremen Spatial Cognition Center, University of Bremen, Germany.

Michael Greenstein, Department of Psychology & Philosophy, Framingham State University, United States.

Zach D. Haga, Center for Applied Brain & Cognitive Sciences at Tufts University, United States.

Lindsay A. Houck, Center for Applied Brain & Cognitive Sciences at Tufts University and Cognitive Science Team, U.S. Army Natick Soldier Research, Development, and Engineering Center, United States.

Andrea Kantrowitz, Department of Art, State University of New York at New Paltz, United States.

Elizabeth J. Marsh, Department of Psychology & Neuroscience, Duke University, United States.

Chiara Meneghetti, Department of General Psychology, University of Padua, Italy.

Jeffrey V. Nickerson, School of Business, Stevens Institute of Technology, United States.

Jane Nisselson, Virtual Beauty, United States.

Ana-Maria Olteţeanu, Bremen Spatial Cognition Center, University of Bremen, Germany.

Francesca Pazzaglia, Department of General Psychology, University of Padua, Italy.

Holly A. Taylor, Center for Applied Brain & Cognitive Sciences at Tufts University and Department of Psychology, Tufts University, United States.

Jasper van de Ven, Bremen Spatial Cognition Center, University of Bremen, Germany.

Jeffrey M. Zacks, Department of Psychological & Brain Sciences and Department of Radiology, Washington University, United States.

1

INSIGHT IN MIND

Holly A. Taylor and Jeffrey M. Zacks

Representations of Barbara Tversky

How do representations in our heads and representations out in the world interact to produce human behavior? This is one of the great questions of contemporary cognitive science and the theme of this book. These questions also reflect the career of Barbara Tversky. The chapters presented here all take aim at cognitive representations in complex cognition. They emerged from a symposium held in Chicago in November, 2015, to honor Barbara Tversky's career. They extend beyond traditional definitions of cognitive science by representing a convergence of psychologists, artists, computer scientists, philosophers, and designers.

This volume is broad in scope, befitting the panoramic range of Barbara Tversky's scientific career. In a recent conversation with one of us, Tversky referred to Yuval Noah Harari's contrast between hunter-gatherers and farmers (Harari, 2014) and characterized herself as an intellectual hunter-gatherer. Indeed, her restless scientific mind has ranged over verbal memory, scenes, events, diagrams, and mental imagery. She has studied college students and children, architects and artists. With her collaborators, she has invented tasks, wrestled behaviors such as wayfinding, drawing, and graphing into the laboratory, and conjured methods to render the dynamic processes of free-range cognition accessible to quantitative analysis. At the same time, within this vast range of topics, she masterfully connects ideas both within and across topics, thus creating order and theory relevant for understanding and predicting complex cognition more generally.

This volume aims to reflect not only Barbara's research but also her theoretical approach, which is unique in its ability to provoke insight and connection. Across the wide range of her work stretches a common focus on the format of representations. Representations can be in one's mind or made explicit on a piece of paper or computer screen in the world. Barbara clearly also recognized that even considering both internal and external representations would be insufficient to understand complex cognition. The mind interacts with the world and it is important to understand fundamentals of how internal and external representations interact.

Another equally important factor to understanding complex cognition emerges in Barbara's research strategy. Just as she examines internal and external representations and their interaction, she takes the real world into the lab and the lab into the real world. In this way, she has been able to explore how cognition actually functions in important real-world situations. In her work and the work she has inspired, people have navigated environments without ever leaving the lab room, creating mental representations from spatial language, maps, and virtual reality. People have also navigated cities of the world and many a college campus. People have been asked to understand or explain diagrams created as experimental stimuli, and Barbara has set out to understand and find commonalities amongst diagrams others have published to explain a concept. She has explored how artists create, how designers design, and how scientists explain.

This wide-ranging curiosity is evident in her collaborators and students—and in the chapters assembled here. We have arranged the book in three sections, moving from representations in the mind to representations in the world to the interaction of mind and world. The chapters themselves reflect research emanating from the lab, "in the wild" of the real world, and conceptualized in the mind of her students, post-docs, and collaborators based on their inspiring interactions with her.

Representations in Mind

Roberto Casati opens with a philosophical grounding of mental representations. Building on the popular distinction between fast/automatic and slow/effortful modes of processing (e.g., Kahneman, 2011), Casati proposes a new mode that offloads some of the slow and effortful components of reasoning onto the environment. The view that human cognition can only be understood in terms of how it is embedded in the environment of things is called "situated cognition," and it is a view with which Tversky has long been engaged. One way to describe Casati's project is to say that he "situates situated cognition." Nancy Franklin and Michael Greenstein take ideas of mental representation to a highly constrained—and high-stakes—setting: legal testimony. Their chapter reviews how memory and storytelling interact during legal testimony. We anticipate that this chapter will be extremely valuable not just to cognitive scientists but also to legal practitioners. In the final chapter of this section, Tad Brunyé, Zach Haga, Lindsay Houck, and Holly Taylor take on another real-world application of memory: finding one's way around. This work, inspired by Taylor's early work with Barbara (Taylor & Tversky, 1992a, 1992b), explores how spatial mental representations impact knowledge of, and interactions with, the environment. All three of these chapters demonstrate the Tverskyan strategy of uncovering fundamental attributes of memory representations by looking at how they engage with complex, naturalistic tasks.

Representations in World

The second section focuses on the representational artifacts that humans create. It opens with a visual essay by Jane Nisselson, abstracting her short film based on Tversky's paper Visualizing Thought (Tversky, 2011). The essay and the film vividly illustrate how the elements of diagrams function. Michel Denis' chapter zooms in on one diagrammatic element, the arrow, wonderfully taking a common artifact that appears simple enough to take for granted and showing the actual cognitive complexity that underlies its functioning. Increasing the complexity of the representational artifacts, Mireille Bétrancourt takes on a particularly contemporary cognitive artifact: animation. It turns out to be surprisingly difficult to design effective animations; Bétrancourt's chapter investigates why this is so and suggests what can be done. Representational artifacts exist for people to use, even if only for the artifact's creator. Jeffrey Nickerson moves from the artifacts themselves to how practitioners really use them, showing how designers use artifacts to reason, to discover, to provoke themselves to new insights. Finally, Francesca Pazzaglia and Chiara Meneghetti show that there are dramatic differences in how people use cognitive artifacts to learn about spatial environments, and that these differences are systematic.

Interaction of Mind and World

We close, of course, with how representations in the mind and in the world interact. Elizabeth Marsh and Kathleen Arnold open this section by considering how using memory cycles back to influence memory representations. Much of Tversky's early research investigated the operations used by children and adults when encoding and remembering verbal materials (e.g., Tversky, 1973; Tversky & Teiffer, 1976). When Marsh and Tversky began collaborating, they focused on how these mechanisms function in the sorts of things that people actually do with verbal memories (Tversky & Marsh, 2000). When people use verbal memories, they interact in interesting ways with the world and come back to influence the mental representation. One important thing people do is tell others about what happened—and as Marsh and Arnold demonstrate, this affects not only the listener but also the memory of the teller. Christian Freksa et al's chapter applies this mind-world interaction analysis to spatial problem solving. Notably, the arrangement of objects in space, whether they be buildings in a city, rooms in a building, or matchsticks on a table, affects spatial problem solving related to those objects. Jeffrey Zacks's chapter applies a similar analysis to event representations, showing how our mental and neural representations of events determine our media—and vice versa. Finally, Andrea Kantrowitz illustrates (literally!) a fine-grained description of how artists go back and forth with their marks, shaping their creative cognition in real time.

The Complete Picture

Many of the authors represented have been committed to the field of cognitive science for long years. Others are recent converts or fellow travelers. We hope that this book will be useful to both old hands and new recruits. For ourselves, we think that the work collected here renders vivid the power, breadth, and creativity of what contemporary cognitive science can be. This only makes sense, given that it was inspired by Barbara Tversky.

Authors' Note

Preparation of this volume and the introductory chapter was a joint and equal effort. Order of authorship was determined by a coin flip.

References

Harari, Y. N. (2014). *Sapiens: A brief history of humankind*. Random House.

Kahneman, D. (2011). *Thinking, fast and slow*. Macmillan.

Taylor, H. A., & Tversky, B. (1992a). Descriptions and depictions of environments. *Memory and Cognition*, 20, 483–496.

Taylor, H. A., & Tversky, B. (1992b). Spatial mental models derived from survey and route descriptions. *Journal of Memory and Language*, 31, 261–282.

Tversky, B. (1973). Encoding processes in recognition and recall. *Cognitive Psychology*, 5, 275–287.

Tversky, B. (2011). Visualizing thought. *Topics in Cognitive Science*, 3(3), 499–535.

Tversky, B., & Marsh, E. (2000). Biased retellings of events yield biased memories. *Cognitive Psychology*, 40, 1–38.

Tversky, B., & Teiffer, E. (1976). Development of strategies for recall and recognition. *Developmental Psychology*, 12, 406–410.

Representations in Mind

2

TWO, THEN FOUR MODES OF FUNCTIONING OF THE MIND

Towards a Unification of "Dual" Theories of Reasoning and Theories of Cognitive Artifacts

Roberto Casati

The main aim of this contribution is to stabilize and generalize the use of the conceptual labels originating from "dual" theories of reasoning, so as to provide a theoretical unification with theories of cognitive artifacts, and to describe in an abstract way the mechanics of cognitive artifacts. Psychological literature has by and large accepted the distinction between two "systems", or – as I shall say – two *modes* of operation of the brain in certain tasks, mainly reasoning and decision-making tasks (Evans 2003, 2012, 2015 for reviews; Evans and Frankish 2009; Kahneman 2011). Mode 1 (M1, for brevity) is an automatic, autonomous, stimulus-driven, fast operating mode that delivers rough but locally acceptable results; M2 is modulated by will and attention, operates slowly and stepwise, intensely uses working memory, and is in general more accurate. I shall take the distinction for granted (with some caveats, in particular I shall argue that we do not need to endorse a substantive view of cognitive *systems*, as opposed to a more neutral talk of *modes*) and argue for an extension of the conceptualization to cover cases discussed in the literature of cognitive artifacts, with the goal of unifying the two fields. I'll first introduce M4, an operating mode that completely outsources the computations typically run by M1 and/or M2 to external artifacts. The M4 mode fully delegates the relevant mental activity – what I shall dub "core" computational tasks – and only makes its user care about the input and output of the computation. Then I'll vindicate the existence of a Mode 3, best understood as occupying an intermediate position between M2 and M4. In the third mode we interact with cognitive artifacts (such as maps, measuring instruments, written text) and this interaction is both essential to performing a certain task (as opposed to what happens in M1 and M2) and is not an instance of wholesale offload (as opposed to M4). Interactions with cognitive artifacts actually display proprietary computations, which give some hints about the architecture of cognition, and about its flexibility. Flexibility in turn creates room for the activity of designers of cognitive artifacts.

In this chapter I investigate the relationships between the M3, M1 and M2, and discuss some demarcation issues: whether M3 activities are a subclass of M2 activities, and whether we should postulate continuity between the four modes. More specifically, I look into some of the proprietary computations of M3 (such as shunting information, bridging cognitive modules, displacing search processes, or restructuring memory search). Other targets are popular metaphors such as the "extended mind" and "the world as external memory" that, by making the M3 look too much like M4, risk missing out on the specific properties of brain-artifact interaction.

M1 vs M2: An Example from Navigation

I grew up in a right-driving country. My parents taught me how to cross the street when I was very young. Over time I refined this practice. I moved from having no idea about crossing streets to being moderately skilled, to being an expert. When you acquire such a skill, the imperatives are to train and to aim at being error free. Over such a long, endless process of learning and perfecting, your responses inevitably become automatic: they involve a sense of your body, a sense of your target, time planning, quick decision making, and gut feelings. For instance, when you cross the street, you automatically expect and almost feel danger first on your left-hand side, then on your right-hand side. Navigation is computation (Hutchins 1995).[1] You assess time in relation to the space of travel, the contingencies of the street scene, your assessment of safe zones, the presence of other actors. You even break the law from time to time. You face, and make, quick decisions if an unexpected car shows up. Crossing the street has become second nature to you; you are *fluent* in crossing the street.

On my first visit to England, a left-driving country, I was back at square one. No more quick assessments, no more quick decisions, no more confidence: I was a child once more. I was warned over and over again about looking "in the other direction" or, if in doubt, "look in all directions". I paid attention to each step I took – an attitude that may even be dangerous in some circumstances, as I would often freeze in the middle of the street. I really planned each crossing in advance, aiming to figure out how to complete it and trying to avoid the prospect of needing a contingency plan (too difficult to figure out, almost impossible to implement). I could almost feel my thinking slowing down. I engaged in hypothetical reasoning, without relying on any past experience of crossing in a similar situation (Evans 2003.) People who have spent enough time in both driving cultures report that they adapt, becoming as skilled as the natives, or at least as fluent as proficient second-language speakers. I clearly did not spend enough time there. Statistics on road accidents in the UK confirm difficulties for right-hand-driving truck drivers (Danton, Kirk, Hill 2009).

There are many entry points to the two-mode metaphor and the reason I choose this one is mainly because it makes the important aspect vivid (if you ever had the experience of the other-driving habit): both a clear phenomenology and obvious behavioral differences are associated with the deployment of each mode in the two different cases respectively. Mode 1 is the way my brain operates when I cross the street in France, and Mode 2 when I do so in Great Britain. In the first case I think fast, in the second I think slow (Kahneman 2011). But what do "fast" and "slow" mean? It is partly a matter of decisions/calculations unfolding in time, but the time differences reflect architectural differences in the organization of thought.

What the Distinction Is and What It Is Not

The psychological literature has accepted a distinction between two "systems", or two modes of operation of the brain in certain tasks, mainly reasoning and decision-making tasks. Mode 1 is described as an automatic, fast and fluent operating mode that delivers rough but locally acceptable results; Mode 2 is described as modulated by will and attention, as operating slowly and stepwise, as taxing memory working, and as delivering comparatively more accurate results in many cases. There are now a number of ways to present the distinction and to relate it to other older distinctions (intuition vs. reason, heuristics and biases vs. logic, encapsulated "Fodorian" modules vs. central processing; see Evans 2003), but let me just say what I take the distinction *not* to be. First and most importantly, it is not a distinction between two *systems*, *strictly considered*, if by "system" we mean something like the visual system or the auditory system, although it may be so interpreted.[2] Kahneman himself insists on the fact that "System one" and "two" are akin to fictional characters. Second, it is not a distinction between two modules (like the vision module for face recognition and the vision module for line orientation). Both M1 and M2 rely on the operation of modules – although in addressing a particular task, they may tap into different modular resources, and deploy them in different ways. Third, and by (almost) the same token, it is not a distinction between having intuitions and not having intuitions. You have and use plenty of intuitions when you operate in M2. Finally, the distinction is not that of innate vs. acquired. Learning does make a difference in operating M2, but on the one hand a lot of learning is present in the myriad systems that operate in M1, and on the other hand – to mention but an example – reading is the result of learning, but competent reading clearly displays M1 automaticity and fluency.

On a positive characterization, M1 and M2 are best construed as *modes of operation* of the brain. This is to say that although based on a host of different competencies, they are better seen at the level of *performance* rather than at the level of *competence*. This is suggested by the literal adverbial reading of Kahneman's distinction between thinking "fast" and "slow". One may thus characterize the approach defended here as an "adverbial" construal of the relevant phenomena, as opposed to a "substantival" construal, typified by talk of "systems".

Generalizing the Distinction

Talk of modes is crucial to a first generalization. There are a large variety of mental engagements over and above verbal problem solving and decision making. As we have seen, navigation is hospitable to the two-modes distinction. Visual analysis is another example: you can just let vision process information automatically, or you may want to intervene on it by directing your overt attention to details (Cavanagh 2004).

Second, in the tasks standardly discussed in the literature on reasoning, the overarching "control concept" of M2 is *enhanced epistemic or pragmatic quality*.[3] This too can be generalized. You want to end up in a zone in which the quality of your representation or of your action is better than the one you would obtain in using the first thing that comes to mind, i.e. in using M1, relative to a certain task such as navigating or finding the solution to a logic problem.

For the present purposes, the main aspect of M2 is that it is a mode of operation in which various M1-like activities are modulated and regimented for performing a certain task such as navigation. When you address a certain problem in M2 mode, you do not trust, or you do not care about, the first thing that comes to mind; and you try to make sure that other, more trustworthy things come to mind. As an illustration, consider the difference between walking down to the cafeteria (something you did so many times that not only does it feel automatic, it is automatic to the point of making you forget that you wanted to go to the news-stand), on the one hand, and on the other hand going from to Place de la Concorde to the Eiffel Tower, unassisted by maps (something that requires you to plan, visualize, double check, and put together bits and pieces of spatial knowledge).

How does M2 work? After a long period of attention to how M1 can interfere with normative reasoning, the literature has produced a number of insights into the functioning of M2 activities. The mechanics of M2 is assumed to be a mix of inhibition and selective attention (Stanovich 2009), but it also includes designing strategies or learning tricks to generate good solutions. For instance, while M1 comes up with biased, perception-generated answers to logic problems such as that presented in the Wason selection task (Wason & Evans 1975), training in logic, or the use of local metacognitive strategies (Houdé et al. 2000) may inhibit the responses that actually come to everyone's minds (including to the logician's mind) and buy time and attention for solving the task, based on principles. Overriding M1 is an interesting design problem for cognition. Probably metacognitive feelings are involved in getting M2 started. For instance, we may detect a conflict and become suspicious; a certain outcome may feel "fishy"; expertise plays a role here.

Once the revision process has started, checking each step of your solution-generating process is a standardly applied strategy. This can only work, of course, if you are able to access the various steps of the computation. Some other strategies are generally available:

- Actually thinking twice, i.e. going over the process one more time: maybe the generated solution differs, or corrects the first one.
- Being systematically suspicious about the result, so that you may want to think twice.
- Learning by heart some partial solutions you are bound to use (for instance, the results of the multiplication table).
- Running the process in reverse, conditionally: if this is what I obtained, then . . .

But, as I said, the actual mechanics of M2 operation are as yet not fully understood (Stanovich 2009). For the present purposes, we can accept that as M2 operation is algorithmic in nature, action can take place at the input, at the computation, or at the output levels. Whether and how this is possible will depend on the type of process you plan to act upon.

Introducing M4

My proposal is to extend the M1/M2 metaphor to include another mode, M3. Actually, I propose to introduce *two* other operation modes, M3 and M4, and in order to exactly understand the import of the introduction of M3 I find it convenient to start from M4. M4 occupies an extreme of the spectrum of ways of operating under examination here. Working in M4 mode means *delegating* the *relevant* mental process to some external *device*, and only caring about the *solution* of the delegated process. M4 really is a proxy mode. For instance, you launch a navigation query ("Destination: Eiffel Tower"), and take directions from a GPS-based navigation device, by listening to the instructions dictated to you over a speaker. "Proceed to the intersection. At the intersection, make a left turn." You do not even look at the map on the display. This is a paradigmatic case of M4 in operation.

Let us spend some time on this particular example of the "talking navigator". First of all, this extreme type of human-machine interaction is only possible in heavily constrained environments, such as the road system, where decisions tend to be discrete, as the options are simplified: you rely on an ontology of networked one-dimensional objects (routes), left turns, right turns, and little else, where the only legal moves are going ahead or turning back – not, say, head for a 37° bearing.

Second, in most cases – and in this one in particular – you do not give up all of your mental life. You are still making decisions, you launch a navigation query, you listen to the verbal instructions and interpret them, you check them against your perception of the environment – conflicts may arise ("was it the first or the second exit at the intersection?") – and you should still be able to solve them. Yet, you are giving up most of the *relevant* mental life, if compared to a more engaged relationship with, say, a paper map. In particular, you are giving up *navigational* computations, that you outsource completely, in the same sense in which you outsource them to a taxi driver in an unknown city (Marconi 2005). In M4 mode you only access/perform at the input and/or the output of the computation, and you do not run any of the relevant computations in your brain.

Let me introduce a distinction between *core* computational tasks and *auxiliary* computational tasks. A core computational task, relative to a given performance, is a task that is essential to obtain the desired result. In navigation, core computational tasks are dead reckoning and determining a direction, for instance. You must be able to compute your local position and you must be able to compute your direction to target, either by finding a heading directly to target, or by laying down a sequence of waypoints and then finding a heading to each of them in sequence. In M4 mode the core computational task of navigation is clearly offloaded. You do not need to know where you are and you do not need to calculate any of the next steps. Still, you must compute something, i.e. how to launch a query (how to interact with the device) and how to interpret the result. Relative to navigation, these are auxiliary computational tasks, as they are not navigational computations. Besides, they are not specific to navigation and are actually shared by many other tasks (e.g. looking up an encyclopedia entry on your PC).

A signature feature of M4 is then that you do not need to access any steps of the core computational task that characterizes the performance. The nature of the performance changes: it is no longer "navigation", but "use of an instrument that finds the route for you."

Introducing M3

I sought to characterize M4 in order to gain a better understanding of M3. M3 is distinct from M2 in that it essentially uses external artifacts, and is however distinct from M4 insofar as, relative to a particular task, delegating to external artifacts is not wholesale – and definitely does not outsource some *core* computational task. As an instance, consider navigating an unknown environment by using a paper map. First, this is clearly different from navigating (M1-like) in a known environment for which your brain has constructed – you can't help it – an effective mental map. The notion of a mental map may itself be controversial (see Madl et al. 2015 for a review), but I assume that information is stored in the brain in map-like, i.e. spatially organized, format. The M1-like features of mental map representation and use, on the other hand, are not controversial. You do not need to consult an external map to find your way back to the elevators in this building: fluency and automaticity are granted here. Second, navigating with a paper map in an unknown environment is clearly different from taking directions from a GPS, or setting an autopilot, which are prime examples of M4 mode of operation. In an unknown environment, assisted by a map, you try in a careful, stepwise fashion to identify landmarks, to interpret and assess affordances, and to keep an explicit mental record of visited places; you are also inhibiting some of your powerful M1 routines (Stanovich 2009). (For instance, your perception-based intuitions suggest proceeding in a straight line, but the map signals an obstacle that is invisible from your current location, and so you take a detour; to take another example, your mental mapping system hooks on to whatever visual features are available in the landscape, including clouds and waves, which make for very ephemeral landmarks, and you have to inhibit or discard this hooking.) You interpret the map propositionally, not just as a depiction of the terrain, but as a set of instructions, of inferential cues, that are subject to distortions of various type (Tversky 1981; Tversky and Lee 1998, 1999; Denis 1997.) All this puts you in an M2-like mode of operation. But the map is essential here; interaction with the map modulates all the other activities you engage in. Map-assisted navigation at its core consists of deliberately thinking of and making a connection between you and your environment, which requires interpreting both the map (to make quick inferences about what to expect around you) and the environment (to assess the presence of landmarks), before trying to connect the one to the other.

A map stores information about a place in a specific, spatially structured way (Casati and Varzi 1999; Rescorla 2009; Kulvicki 2015.) For our present purposes, the important point is that it records information about places that you may or may not have visited, and in particular about places that you may not be able to perceptually access in the present moment. A map used for navigation should be capable of supporting certain specific queries. By asking the appropriate questions of the map, you can for instance determine that you are in a certain location, and that you should move in a certain direction if you want to reach your final destination. In order to perform these functions in navigation, a map needs *orienting.* Orienting can be explicit (as when you are able to have a meridian's representation on the map coincide with the represented meridian), or implicit (you master a set of rules for transforming each direction you establish on the map on a direction in the represented region; no matter the actual position of the map relative to you and your environment). In both cases you are establishing *physical relationships* between the map and the environment in which the map is situated, and this requires a certain set of skills on your part, which you must learn. Navigation is computation: you should be able to compute, on an oriented map, your current location, by finding the appropriate correspondences between represented landmarks and their representations (for instance, the bearings of two landmarks intersect on the map in a point that approximates your current location).

Are you "delegating" to a map, in the sense in which you were delegating to the talking navigator? Certainly not. When you interact with a map you engage in a specific mode of operation. To recap, you have to be able to ask questions that the map can provide answers to, and be able to find the answers on the map, in order to make decisions; you actively engage with the map, and this engagement lets you generate the mental representations that you need to make your decisions and start your motor routines. You are core-computing, on the map and with the map, position and route yourself. This way of interacting I take as definitional of M3.

The Mechanics of M3

M3 involves external cognitive, epistemic artifacts essentially, and we have to consider now the features of these artifacts. The first key distinction I have in mind here is that between representational vs. non-representational cognitive artifacts.[4]

Some artifacts are not cognitive: hammers, nails and scissors belong in this family. Although you have to cognitively engage with them in order to use them, and although they may have a cognitive effect on you, their function is not to impinge on your cognitive states; whatever they may do to cognition, they don't do it by design. On the other hand, epistemic artifacts, by design, are meant to have an effect on your cognitive states. For instance, they may have been designed to improve the quality of your cognitive states. How do they do this? The main sub-distinction here is between representational and non-representational artifacts. A pair of glasses with suitable lens correction will make you see better. If you have to read a road sign, you end up in a better cognitive situation relative to not wearing them. Glasses intervene on the incoming information by structuring it. However, glasses do not represent anything; they are not about anything. Wearing glasses does modify your mental representation – your visual image of the scene in front of you – but glasses (and their cognates mirrors, microscopes, telescopes, periscopes) are not like written words, maps or painted images, which are items that have representational content (over and above structuring the information that gets to your senses).

In the case of cognitive artifacts, the pragmatics condition the recruitment of cognitive functions and strategies and constrain the design of the artifact. Given the variety of uses, we observe some quite specific and at times interconnected ways cognitive artifacts interact with the brain, i.e. types of cognitive mechanics:

> freeing up working memory, storing information, making inferences visually available and available for inspection, priming action, bridging different modules, shunting information, providing hyper-stimuli to a module, making orienting and navigation possible, indexing, displacing complex searches to automatic brain routines.

This is, of course, an open-ended list. Here is where a key theoretical tool of the account I would like to propose finds its place. In all those cases we can point to some specific *cognitive advantage* linked to the adoption of a given artifact. One of the main tasks for psychology is then to find operationalizations of the advantages (e.g. in what is it easier to use a route display rather than a map display on an embarked navigation system? What makes certain types of diagrams better suited than others in order to extract inferences? (Tversky et al. 2000, Heiser & Tversky 2006 on arrows; Tversky 1995); Why are certain items in the multiplication table easier to remember than others?) And the main task for designers is to assess and exploit the trade-offs between cognitive advantages.

Let me give a few examples of quite specific interactions between cognitive artifacts and the brain along the lines I just drew.

Displacement. Suppose you are looking for the trains leaving between 9 and 11. You are presented with four timetables, each presenting the same facts, but in a different graphic style: Mixed Bag, Order, Highlight, Chunk. You shall notice a difference in retrieving the desired set of times.

(Mixed Bag):	10:20, 8:30, 11:40, 13:30, 9:25, 10:35, 16:40, 8:15
(Order):	8:15, 8:30, 9:25, 10:20, 10:35, 11:40, 13:30, 16:40
(Highlight):	**10:20**, 8:30, 11:40, 13:30, **9:25**, **10:35**, 16:40, 8:15
(Chunk):	8:15 8:30, 9:25, 10:20,10:35, 11:40, 13:30, 16:40

Larkin and Simon (1987) would claim that the four timetables are informationally equivalent but computationally different. Tversky et al. (2007) distinguish here between the Principle of Congruence ("the structure and content of a visualization should correspond to the structure and content of the desired mental representation") and the Principle of Apprehension ("the structure and content of a visualization should be readily and accurately perceived and comprehended", p. 56). But these overarching distinctions do not tell us about the mechanics yet. What are the computational differences that matter for Apprehension? Why does it take longer to search in (Mixed Bag) than in any other style? If you think about your behavior in front of (Mixed Bag), unassisted by paper and pencil, you'll notice that you have to pay attention sequentially to each item in the series, you have to explicitly mentally label it as a positive or a negative, you have to store positives in memory, and possibly revisit the series for double checking. That means that both targets and distractors are processed and re-processed. You must remember what the target's content is, or where an encountered target was, in case you forgot the content. Both spatial and short-term memory are actively engaged and put under pressure, and paying attention to each item clearly locates your activity in an M2-like area. In the other three formats, powerful and fast computational routines of the visual system take up part of this chore, kicking it to M1 modes of operation. The visual system loves ordered series, color differences and chunks, and can instantly zero in on the positives as they have been singled out in a visual-system-friendly format (Healey and Enns 2012). Displacement of cognitive tasks is one of the key mechanisms of cognitive artifacts. The design implicit in (Order), (Highlight), and (Chunk) is an orchestration of M1-operating routines, that are more effective than the M2 work done when searching in (Mixed Bag).

Displacement is interesting in two other respects. First, it does not require learning, and thus is extremely effective, giving designers of cognitive artifacts great leverage. Second, moving computations to one part of the brain to another keeps the computations within the boundaries of the brain. The non-necessity of learning distinguishes Displacement from other shifts of cognitive load within