



Published in final edited form as:

J Exp Psychol Gen. 2014 June ; 143(3): 1314–1330. doi:10.1037/a0035542.

MAGELLAN: A cognitive map-based model of human wayfinding

Jeremy R. Manning^{1,2}, Timothy F. Lew³, Ningcheng Li⁴, Robert Sekuler⁵, and Michael J. Kahana³

¹Princeton Neuroscience Institute, Princeton University

²Department of Computer Science, Princeton University

³Department of Psychology, University of Pennsylvania

⁴Department of Bioengineering, University of Pennsylvania

⁵Volen Center for Complex Systems, Brandeis University

Abstract

In an unfamiliar environment, searching for and navigating to a target requires that spatial information be acquired, stored, processed, and retrieved. In a study encompassing all of these processes, participants acted as taxicab drivers who learned to pick up and deliver passengers in a series of small virtual towns. We used data from these experiments to refine and validate MAGELLAN, a cognitive map-based model of spatial learning and wayfinding. MAGELLAN accounts for the shapes of participants' spatial learning curves, which measure their experience-based improvement in navigational efficiency in unfamiliar environments. The model also predicts the ease (or difficulty) with which different environments are learned and, within a given environment, which landmarks will be easy (or difficult) to localize from memory. Using just two free parameters, MAGELLAN provides a useful account of how participants' cognitive maps evolve over time with experience, and how participants use the information stored in their cognitive maps to navigate and explore efficiently.

Keywords

spatial navigation; cognitive map; allocentric; spatial memory; virtual reality

Introduction

“...[S]omehow you're suddenly lost, driving with mounting wretchedness and confusion, fully aware that the clock is ticking and you're going to be late for your child's soccer game or that important dinner party. Most of the time there is nothing for it but to ask for directions from a passerby or stop at a gas station or 7-Eleven. Even then you are likely to slightly misunderstand what you're told so that you need to repeat the same shameful inquiries...”

—John Edward Huth, *The Lost Art of Finding Our Way*

Anyone who has become lost while searching for a destination in an unfamiliar environment intuitively appreciates the cognitive challenges presented by this all-too-common situation. These challenges include the acquisition, storage, processing, and retrieval of spatial information, all of which are carried out while one moves through the environment. How do our brains carry out these complex tasks? For the past half-century, neuroscientists have hypothesized that, by modifying their firing rates at specific locations within an environment, specialized *place cells* originally discovered in the rodent hippocampus (O'Keefe & Dostrovsky, 1971) form the basis of an allocentric cognitive map (i.e., a map referenced to fixed navigationally relevant objects (landmarks) rather than egocentric coordinates; Tolman, 1948). The discovery of entorhinal *grid cells*, which respond at the vertices of triangular lattices tiling an environment and serve as a major input to place cells (Moser, Kropff, & Moser, 2008), has begun to further elucidate the neural basis of the rodent cognitive map. Over the past decade, advances in neural recording techniques (e.g. Fried et al., 1999) and virtual reality technology have facilitated the discovery of similar populations of place cells (Ekstrom et al., 2003) and grid cells (Jacobs et al., 2013) in the human brain. This suggests that human navigation and rodent navigation may rely on similar allocentric representations of navigated environments. To test this hypothesis, we developed a computational model named MAGELLAN,¹ which assumes that spatial information is stored and retrieved from an allocentric cognitive map.

Our primary research focus is to understand how humans learn to wayfind efficiently in unfamiliar environments. Previous work has ratified the common sense intuition that when we first encounter an unfamiliar environment we initially have difficulty finding our way around, but we learn to navigate more efficiently with experience (e.g. Newman et al., 2007). We sought to identify the principal dimensions of information acquisition, storage, processing, and retrieval that support experience-based improvements in navigation. Specifically, we sought to model the dynamic process by which people learn to navigate in an unfamiliar environment, from acquiring information about the environment, to storing it (and forgetting information over time), to using the stored information to navigate more efficiently. Our model also explains how a navigator can carry out efficient search for an unknown target, using existing spatial knowledge.

To study how people learn to navigate efficiently in unfamiliar environments, we asked human participants to navigate computer-generated virtual towns that they had not seen previously. Participants in our experiments played the role of taxicab drivers who had to pick up and deliver passengers to specific locations in the towns. We used the MAGELLAN model to estimate the moment-by-moment contents of participants' cognitive maps by taking into account participants' past experiences with specific objects in the environment. Given the moment-by-moment cognitive map estimates, we used MAGELLAN to estimate the subsequent paths that a participant would take in navigating in the environment. We then tested the model by comparing participants' data with the model's predictions.

In the next section, we present the key details of the model that guided our empirical work. We then describe an experiment that incorporates a navigational task, and explain how

¹Our model's eponym is the Portuguese explorer Ferdinand Magellan, commander of the first ship to circumnavigate the globe.

MAGELLAN's parameters can be adjusted such that the model's behavior matches that of participants. Next, we use the model to generate a series of new environments that MAGELLAN predicts should constitute a range of challenges for potential navigators. Our second experiment tests human navigators in a subset of these model-generated environments. Our data confirm MAGELLAN's predictions concerning (i) the relative overall difficulties people have in navigating the different environments, (ii) the relative rates at which people learn to navigate the environments efficiently, and (iii) the memory strengths associated with specific landmarks.

The MAGELLAN model

The MAGELLAN model comprises three modules: a *vision* module that processes visual information acquired during navigation, a *cognitive map* module that stores spatial information about the environment, and a *route generation* module that uses the information stored in the cognitive map to navigate towards the current target (Fig. 1). The model's behavior can be tuned by adjusting two scalar parameters, V and M . Each is an aggregate that lumps together multiple components into a single parameter. The vision parameter, V , controls the efficacy of the model's vision module by defining the proportion of the computer screen that a structure² must occupy in order to be seen by the model and entered into memory. The memory parameter, M , controls the efficacy of the cognitive map module by defining the time that spatial information stored in the cognitive map remains viable. Using just these two free parameters (holding the efficacy of the route generation module fixed), MAGELLAN successfully captures a wide variety of navigation behaviors ranging from efficient directed search to random walks. The interactions between MAGELLAN's three modules provide an account of how a navigator acquires spatial information, learns the spatial layout of a new environment, and navigates to known and unknown targets in the environment. In doing so, MAGELLAN explicitly estimates the moment-by-moment states of a navigator's cognitive map based on what they have seen and done, and predicts the future path a navigator will take towards a target given the spatial structure of the environment and the body of spatial knowledge represented in the navigator's (estimated) cognitive map.

An ideal navigator as a benchmark

Everyday experience too often reminds us of our imperfections as navigators. For example, when we find ourselves in an unfamiliar environment, we get lost, we fail to notice potentially important features, and sometimes we forget some of the things that we did notice. In order to characterize how imperfect human navigators would perform in unfamiliar virtual environments, we compared their performance to the performance of an ideal navigator. By definition an ideal navigator, armed with knowledge of its current position in an environment and the position of an intended destination (target), will generate an *optimal* path from its current position to the target. In the context of our taxicab task, an optimal path is a delivery path whose length is equal to the minimal achievable path

²Following Han, Byrne, Kahana, and Becker (2012), we use the term *structure* to reference objects in the environment, independent of participants' experiences with those objects. Structures become *landmarks* when they become navigationally relevant, for example when they are added to spatial memory.

distance, taking into account the need to detour around impenetrable obstacles. Note that an ideal navigator is not prescient: it cannot generate optimal paths to targets it has not yet seen. For targets that have not yet been seen, the ideal navigator employs an efficient search algorithm, quickly searching for a target in the nearest unknown sections of the environment. Once the target is seen and added to the ideal navigator's cognitive map, the navigator's route-generation module produces the minimal achievable path from the navigator's current location to the target.

The ideal navigator incorporates all visible structures into its cognitive map, and remembers those landmarks indefinitely and with perfect fidelity. As a result, the ideal navigator's spatial knowledge, which is stored in its cognitive map, provides an upper bound on the spatial knowledge an actual human navigator could have about the environment's layout. Because the ideal navigator generates optimal paths to known targets and searches efficiently for unknown targets, the ideal navigator's paths provide an upper bound on the efficiency with which a human navigator could navigate to those same targets. With these two upper bounds as benchmarks, the performance of human navigators can be compared to the performance of the ideal navigator. This comparison can yield important insights into sources of errors that human participants make, opening a window onto human navigators' perceptual and cognitive limitations.

In our explorations, we systematically altered MAGELLAN's vision parameter, V , and its memory parameter, M . Variations in the V parameter govern which structures in the environment are added to the model's spatial memory; variations in M control how long landmarks in that spatial memory remain viable. When V and M are set appropriately (as described below), MAGELLAN operates as an ideal navigator with perfect vision and spatial memory. By adjusting the value of the V and M parameters, we can systematically degrade MAGELLAN's vision and cognitive map modules, respectively, to more closely match participants' behaviors.

Vision parameter— V . MAGELLAN's vision parameter, V , defines a threshold fraction of the visual field displayed on the computer screen, $0 \leq V \leq 1$. This is the fraction of the display that a structure must occupy in order to be perceived, and thus added to the model's cognitive map. The V parameter is intended to account for failures to encode some structure displayed on screen. Such failures could occur because the structures occupied too small a fraction of the screen and/or because the participant brought insufficient attention to bear on the given structure as they navigated the virtual town. Intuitively, the V parameter may be thought of as representing vision's spatiotemporal limitations (e.g. Geisler, 1989).

When $V = 0$, any structure displayed on the screen is added to memory, and so MAGELLAN's vision behaves as an ideal navigator's would. As V increases, MAGELLAN's vision grows less effective, and the cognitive map and route generation modules are forced to operate with incomplete information. Finally, with $V = 1$, MAGELLAN's vision module can no longer guide navigation, so the now sightless model must rely entirely on blind search to bring itself to a target.³ During blind search, as MAGELLAN drives past each landmark, the corresponding location in the cognitive map is filled in. MAGELLAN then tends towards unexplored blocks that are nearby its location,

working outward towards unexplored blocks that are further away. The “working outward” is driven by our implementation of MAGELLAN’s exploration algorithm (described below), which directs the model to continually move towards the nearest unknown block. This search strategy is reminiscent of *spiral search*, an optimal blind search algorithm employed by desert ants (Müller & Wehner, 1994).

Memory parameter— M . MAGELLAN’s memory parameter, M , is intended to account for forgetting over time. This parameter governs the number of modeled steps after which a newly acquired memory has degraded to the point where it is no longer viable. The motivations for this definition of the parameter are (i) the observation that participants in our task sometimes generate sub-optimal paths to familiar targets (Fig. 3), and (ii) evidence that in other tasks, a navigator’s memory for path details degrades with path length (Lyon, Gunzelmann, & Gluck, 2008).

When $M \rightarrow \infty$, landmarks that are added to the model’s cognitive map remain there permanently. Simulations of the ideal navigator incorporate $M \rightarrow \infty$. In practice, MAGELLAN’s memory behaves identically to the ideal navigator’s as long as M is greater than or equal to the total number of steps taken during a given simulation. When $M = 0$, new information is not at all retained in spatial memory. As a result, when $M = 0$, with each iteration of the simulation, the model takes a step in a randomly chosen direction, which reduces the navigator’s path to a random walk. Note that setting $M = 0$ (i.e., eliminating the model’s ability to remember) has a different effect than setting $V = 1$ (i.e., eliminating the model’s ability to see). Whereas setting $M = 0$ eliminates purposeful navigation entirely (regardless of V), setting $V = 1$ still allows for blind search, provided that $M > 0$.

Implementation

As with most computational models, implementing MAGELLAN requires that we define the set of general experimental paradigms to which the model will apply. In Experiments 1 and 2, participants function as taxicab drivers who pick up passengers, delivering each to a requested destination. While the virtual towns in these experiments comprise regular grids, we note that the MAGELLAN model may be easily extended to more complex or irregular environments by modifying the route generation module appropriately.

We next provide details about Experiment 1, and describe how we implemented the MAGELLAN model to perform the same task that human participants did. After fitting the vision parameter V and memory parameter M to account for the participants’ performance in Experiment 1, we used the fitted model to design Experiment 2, which tests several detailed predictions of the model.

³One set of circumstances required a modification of this application of the V parameter to be altered: If MAGELLAN was both immediately adjacent to a structure and facing towards the structure, then that structure was automatically added to MAGELLAN’s spatial memory – regardless of how much of the screen the structure occupied. This modification kept the model from becoming “stuck” when the V parameter was large, and the model’s vision was so impaired that it could not see a structure that was right in front of it. In this way, the model retained the ability to perform blind search relying on its spatial memory, even if its vision module could not provide useful information.

Experiment 1

Several close variants of our spatial navigation task, known as *Yellow Cab*, have been used in previously reported behavioral and electrophysiological investigations of human navigation (Caplan et al., 2003; Ekstrom et al., 2003, 2005; Newman et al., 2007; Jacobs, Kahana, Ekstrom, & Fried, 2007; Manning, Jacobs, Fried, & Kahana, 2009; Jacobs, Kahana, Ekstrom, Mollison, & Fried, 2010; Watrous, Fried, & Ekstrom, 2011). These previous studies employed relatively small environments, whose layouts participants learned very quickly (e.g. Newman et al., 2007). As we were particularly interested in studying the spatial learning process, we attempted to retard participants' achievements of asymptotic performance by increasing the size and visual complexity of the environments they encountered.

Methods

We tested 21 participants; 14 at the University of Pennsylvania and 7 at Brandeis University. All participants (12 male, 9 female) received monetary compensation, with a performance-based bonus.

Participants learned to navigate virtual towns each laid out on a 6×6 orthogonal grid of streets. A single uniquely textured structure was centered on each block. These structures comprised two categories: thirty-one multistory office *buildings* and five one-story *stores*. The photorealistic textures mapped onto the sides of each structure (Fig. 2A) were drawn from an image pool used in several previous navigation studies (Jacobs et al., 2007; Manning et al., 2009; Jacobs et al., 2010; Watrous et al., 2011). Participants were instructed to pick up passengers and deliver them to specific stores, using a gamepad to control their movement through the environment. Prior to the first testing session, participants practiced using the gamepad controller to navigate to stores in a simplified practice environment. The rates of movement (up to 1 virtual block/second) and turning (up to 13° /second) were determined by the pressure participants applied to the gamepad. Participants were able to alter their heading, by turning in place, without altering their location within the virtual environments. The participants' views from within the environments were displayed on a 27 inch monitor (1024×768 resolution) located at eye level approximately 2 feet from the participants' heads. The rendering engine (Geller, Schleifer, Sederberg, Jacobs, & Kahana, 2007) used perspective projections to give the appearance of depth on the two-dimensional viewing surface of the display, and simulated an infinite depth of field such that all objects displayed on screen were in perfect focus.

The office buildings in the environments varied in width, height, and visual appearance, and were surrounded by sidewalks which participants were not permitted to drive on. While office buildings were never the target of a taxicab passenger delivery, they served two main functions. First, the buildings were intended to provide a sense of spatial context by establishing a set of visual features unique to each spatial location within the environment. Second, office buildings were taller and wider than the stores to which passenger deliveries were made, and thus often blocked the stores from view. This meant that the locations of visually occluded stores had to be learned and remembered if participants were to deliver their passengers efficiently.

The stores, which were all sized identically, were surrounded by pavement that participants could drive on. Stores were distributed pseudo-randomly throughout the town, subject to all of the following constraints: (1) stores could not occupy all four corners of the town, (2) stores could not occupy any two adjacent blocks, (3) stores could not be a “knight’s move”⁴ away from one another, and (4) stores that shared the same North-South or East-West alignment were separated by at least two blocks along the orthogonal axis. These constraints were intended to spread out the locations of the stores and to prevent stores from lying within viewing distance of one another. The outer boundary of the town was marked by a texture-mapped stone wall. A cloudy sky was visible beyond the outer wall.

When a participant first began to navigate in an environment, a virtual passenger was placed a short distance away, in plain view directly ahead of the participant. This placement was meant to reduce the variability in the amount of exploration that different participants would have to do prior to picking up their first passenger. Note that reducing between-participant variability in exploration would control participants’ familiarity with that environment. However, on subsequent deliveries the participant had to search for the next passenger. Passengers (except the first) were scattered pseudo-randomly throughout the environment, subject to the constraint that no passenger could be located in the direct line of sight from the preceding store (i.e., the store to which the preceding passenger had just been delivered). When a passenger was picked up (by driving up to them), a text screen gave the participant the name of the next target store. To deliver a passenger, the participant had to drive directly up to the side of the appropriate store. Over the course of three 30-minute sessions separated by at least one day, participants delivered three passengers to each of the five stores, in each of six unique VR towns (two towns per session). The deliveries were block-randomized such that all stores were chosen as targets before any stores were repeated.

The upper left corner of the screen continuously displayed a short text of the instructions that were currently in effect (e.g. “Find a passenger” or “Find the Coffee Store”). Figure 2A provides a screen capture showing a participant’s eye-view of the environment during a typical search for a passenger. To motivate participants, each one began with 300 points. Points were gradually lost as time was spent driving through the environment, and each successful delivery earned participants 50 additional points. We instructed participants to maximize their total score by learning the layouts of the environments.

Results

Over successive deliveries, participants learned to find more efficient paths from the locations at which randomly placed passengers were picked up to the locations of the stores to which passengers asked to be delivered. To measure this improvement in navigation efficiency over successive deliveries, for each delivery we subtracted the *city block distance* ($x + y$) between each passenger’s pickup location and delivery target from the actual distance a participant traveled between the pickup and delivery. Because the environments were laid out on regular grids, the city block distance between the passenger and destination provides an approximation of the minimum-distance path a navigator could have taken to

⁴A knight’s move is defined as a two block step in one direction (e.g. North) followed by a one block step in an orthogonal direction (e.g. East).

deliver the passenger. We therefore refer to the difference between the city block distance and a participant's actual path length as *excess path distance*. As Figure 3A shows, excess path distance decreased rapidly over successive deliveries. This rapid learning may have been partially facilitated by the geometrically simple layouts of the environments participants navigated in, although we note that this does not, by itself, diminish the environments' realism (for example, many urban areas have similar grid-based geometries). Having established that participants showed rapid learning and near-perfect performance after just a few deliveries, we implemented an ideal navigator version of the MAGELLAN model to assess the degree to which the participants fell short of perfection. We next describe how we implemented MAGELLAN's route generation module to operate in *Yellow Cab* environments.

The route generation algorithm—From the pickup of a passenger to the delivery of that passenger to his/her desired target store, MAGELLAN's route generation module operates via a pair of looping instructions:

1. **Find a goal.** If the location of the passenger's desired destination store is already in spatial memory, set the goal to that destination. Otherwise, set the goal to the closest block not yet associated with a landmark.⁵ If there are multiple equidistant possible goals (after accounting for environmental obstacles), select one at random.
2. **Take a step towards the goal.** If there are multiple equally efficient directions in which to step, select one at random.

A single *step* by MAGELLAN is defined as a move from one intersection in the town's roadways to an adjacent intersection. Consequently, the model always moves in the town along one of the cardinal directions (i.e., multiples of 90°). This simplification allows us to focus on path distance—rather than the exact *shape* of a particular delivery path—as an index of spatial knowledge of the environment. Constraining the model's movement in this way ensures that, unlike human participants, the lengths of the model's navigated paths will be lower bounded by the city block distances from the passengers to their requested destinations. By contrast, a human participant's delivery path could cut corners at headings other than multiples of 90°. Thus, the lower bound of participants' path lengths fell somewhere between the city block distances and the Euclidean distances between each passenger and destination. In order to make fair performance comparisons between participants and the model, we transformed the participants' actual paths to approximate the paths that would be taken by the model by assigning each point along participants' paths to the nearest intersection and connecting the intersections. We used these transformed paths when ever we directly compared participants' performance to MAGELLAN's performance (Figs. 4 and 5).

When MAGELLAN first enters an unfamiliar environment, its spatial memory is a *tabula rasa*. As it moves through the town, the model uses its vision module to note structures that are visible from its current location and heading (assuming the same 60° field of view

⁵Recall that each block in the environment contains a structure—either an office building or a store. Passengers are delivered only to stores.

available to the participants). Each visible office building and store is added at the appropriate location to the model's cognitive map of the town. The cognitive map includes the spatial location and a unique identifier for each landmark that was visible, and fills up as more office buildings and stores are seen.

When searching for a goal store that is not represented in the model's current cognitive map, the model minimizes effort per delivery (Robles-De-La-Torre & Sekuler, 2004), exploring first the most proximate location(s) at which office buildings or stores have not yet been seen. An indirect consequence of this goal-seeking strategy is that previously traversed areas of the town will tend to be avoided while searching for the goal. This strategy is analogous to *inhibition of return*, a phenomenon that has been observed in visual search and other tasks (Klein, 2000; Bennett & Pratt, 2001). Inhibition of return manifests itself through an increased tendency to orient or attend to novel locations, which could promote efficiency in exploring unfamiliar environments or, for example, when foraging for food.

Ideal navigator vs. Experiment 1's human participants—Figure 4 compares the performance of the ideal navigator implementation of MAGELLAN (i.e. with $V = 0$, $M \rightarrow \infty$) to that of participants in Experiment 1. Our goal in this comparison was to determine whether an ideal navigator could predict participants' excess path length on a given trial, taking into account participants' past experience prior to that trial. Therefore, at the start of each delivery path we re-populated the ideal navigator's memory to contain exactly the set of structures that had been visible to the participant, from the start of the participant's session until the start of the delivery being simulated. Because the ideal navigator's memory is perfect, any structure displayed on the participant's screen before the current passenger was picked up was available in the ideal navigator's cognitive map.

We found that the ideal navigator navigated to targets far more efficiently than human participants did (Fig. 4). The magnitude of this discrepancy suggests that one or more of the ideal navigator's parameters substantially overstates the quality of the corresponding function in human navigators. To identify which functional module(s) had been overstated in implementing the ideal navigator, we systematically degraded two of the ideal navigator's three modules (see Fig. 1). As our virtual environments had relatively simple layouts, we assumed that human navigators' abilities to generate efficient paths—if they knew their own location and the location of a target store—were unlikely to differ appreciably from the ideal navigator's ability to do the same. Therefore we preserved the route generation module in its non-degraded form, while degrading the vision and cognitive map modules to assess the degree to which imperfect vision (data acquisition) and/or spatial memory (storage and retrieval) could account for the performance discrepancy between the ideal navigator and our participants.

Degrading MAGELLAN's vision and memory to account for experimental data

—We sought to determine how much MAGELLAN's vision parameter (V) and memory parameter (M) had to be degraded to bring the ideal navigator's performance into line with human navigators. This strategy for treating the components of an ideal, theoretical benchmark system resembles that of Geisler (1989) and others in the sensory domain, and of Robles-De-La-Torre and Sekuler (2004) in the domain of precision motor control.

To optimize MAGELLAN's account of participants' empirical performance, we minimized the root mean squared deviation (RMSD) between the mean *observed* learning curve and the mean learning curve *predicted* by MAGELLAN. Given that the model has only two free parameters, we were able to find the best fitting values with a grid search of the two-dimensional parameter space, varying the vision parameter, V , in increments of 0.01 (equivalent to 1% of the screen) and varying the memory parameter, M , in increments of 1 *step*. The best-fitting set of parameters were $(V, M) = (0.08, 32)$, with an RMSD of 0.652 blocks. Note that with $V = 0.08$, the model's effective field of view is approximately 5° . The rectangles in the lower right of Figure 2A denote 1%, 8%, and 16% of the screen. Figure 5 displays the corresponding predicted and observed spatial learning curves. Recall that MAGELLAN moves one block with each time step, which means that the model's predictions were accurate to within an average of less than one modeled step per delivery. As participants took an average of 11.45 one-block steps between successive deliveries, the best fitting value of the memory parameter ($M = 32$) means that MAGELLAN's memories remained viable for 2.79 deliveries (on average) from the time the memories were initially acquired.

The quality of fit to the mean learning curve from Experiment 1 shows that, with appropriately degraded vision and memory modules, MAGELLAN's performance matches the mean performance of human navigators. However, this does not guarantee that the model can also account for variability in difficulty across environments, or even across trials within a given environment. To the extent that MAGELLAN's behavior mirrors that of human participants, variability in MAGELLAN's performance in unfamiliar environments should predict human navigators' performances in those same environments. Experiment 2 tests MAGELLAN's ability to correctly identify environments that humans will find easy or difficult to learn to navigate efficiently.

Experiment 2

The navigational challenge posed by delivering passengers within an environment in Experiment 1 depends on a number of factors whose interactions cannot be predicted from a simple model-free analysis of the task. These factors include (among others): (1) the placement of stores and buildings within the town, (2) the sizes of the buildings, (3) the placement of passengers, and (4) the sequence of requested destinations. In addition, the routes a participant took on previous deliveries influences the challenge posed by the current delivery, as previous experience determines what information is available to the participant. Experiment 2 was designed to systematically test MAGELLAN's ability to predict the difficulty of learning to navigate efficiently in various environments, based on variation in these factors.

Methods

One hundred and four individuals (37 male, 67 female) at the University of Pennsylvania participated in the experiment for monetary compensation, with a performance-based bonus. As in Experiment 1, participants learned to navigate to stores in 6×6 towns, each containing 31 office buildings and 5 stores. Also as in Experiment 1, participants practiced using a gamepad controller to navigate to stores in a simplified practice environment prior to

their first testing session. The rates of movement (up to 1 virtual block/second) and turning (up to 13°/second) were determined by the pressure participants applied to the gamepad. Participants were able to alter their heading, by turning in place, without altering their location within the virtual environments. The participants' views from within the environments were displayed on a 27 inch monitor (1680 × 1050 resolution) located at eye level approximately 2 feet from the participants' heads. The rendering engine (Solway, Miller, & Kahana, 2013) used perspective projections to give the appearance of depth on the two-dimensional viewing surface of the display, and simulated an infinite depth of field such that all objects displayed on screen were in perfect focus.

Participants were tested in eight unique virtual towns over the course of two 1-hour testing sessions separated by at least 1 day. Whereas the environments used in Experiment 1 were generated randomly and independently for each participant, all Experiment 2 participants encountered the same eight environments (but in different orders). From a set of 500 randomly generated environments, we used the MAGELLAN model to select eight environments varying systematically in predicted difficulty.

To generate a library of 500 towns, we first generated a large pool of office building textures. Each full building texture comprised a bottom texture, containing a door and a row of windows on a brick background; and a top texture, containing a row of windows on a brick background. We used a total of four brick textures, five window styles, four door styles, and three door placement options, yielding 240 possible bottom textures and 20 possible top textures (Fig. 6A). Each top texture could be combined with each bottom texture, producing 4,800 possible combinations. We varied building heights by concatenating either two or five copies of the top texture onto the bottom texture. This generated buildings that were either three or six stories high. We also varied the widths of the buildings by either cropping the full building textures from the sides, or leaving the textures un-cropped. In total we generated 19,200 unique building textures. To make the buildings even more visually distinct, we shifted the hues of each building image (preserving the covariance structure of the pixel intensities), as shown in Figure 6B.

To each randomly generated town, we assigned 31 buildings, selected without replacement from the full set of 19,200 textures. We also assigned five unique stores to each town (Fig. 6C). We required that each store only appear once across all environments a participant encountered. To enforce this constraint, after selecting the final set of eight environments that would be used in Experiment 2, we re-assigned each town's store labels (holding the delivery locations and orders fixed) to ensure that any store would appear only once in the full set of eight environments. MAGELLAN is sensitive to structures' sizes (since this determines which structures are visible to the model), but in our implementation we made the simplifying assumption that the specific textures and store names were task-irrelevant.

Unlike in Experiment 1's towns, store placements in Experiment 2 towns were unconstrained. The constraints on store placements in Experiment 1 were intended to produce towns with spread-out stores and of approximately uniform difficulty. By contrast, in Experiment 2 our goal was to generate towns with a wide variety of layouts. Each of the randomized layouts was also assigned a delivery sequence. As in Experiment 1, these

delivery sequences were block-randomized, with three deliveries assigned to each of the five stores. Thus, each *environment* comprised a combination of a specific town layout and delivery sequence. As shown in Figure 7, removing the store placement and delivery order constraints in Experiment 2 resulted in substantially more variability across environments in both the mean distance between stores ($F(125, 499) = 0.517, p < 10^{-4}$) and distance between successive deliveries ($F(125, 499) = 0.642, p = 0.003$).

Whereas Experiment 1's participants had to search for randomly placed passengers before delivering them to their desired stores, in Experiment 2 we eliminated passenger search. To speed up testing time and to eliminate a less important phase of each trial, after each delivery we simply presented a text screen that specified the next target store. Critically, this change also allowed us to simulate each sequence of deliveries in each of the 500 environments without relying on any behavioral data. This was useful, as our implementation of MAGELLAN has no means of searching for non-stationary targets like the passengers in Experiment 1. We used the best-fitting values of the V and M parameters derived from the Experiment 1 data to simulate 50 runs of MAGELLAN through each environment. We then used MAGELLAN's behavior in these simulations to compute a *predicted difficulty* score for each environment by computing the mean sum of simulated excess path distances for the full set of deliveries across the 50 simulated sessions. Because MAGELLAN explicitly estimates which structures are visible onscreen at each point during the experiment, and which landmarks have remained viable in the cognitive map during each delivery, examining which properties of the environments contribute to predicted difficulty can yield interesting insights (Fig. 8). For example, we found that MAGELLAN predicted that environments with more widely dispersed stores (Fig. 8B) would be more difficult to navigate efficiently, as would environments that required participants to travel further between successive target stores (Fig. 8C). This latter result held even after taking into account the average distances between the stores (Fig. 8D).

Based on the rankings of expected difficulty for the 500 randomly generated environments, we selected eight environments that represented four distinct levels of expected difficulty. The environments we designate as A and E represented the lowest expected difficulty, environments B and F were ranked at the 33rd percentile, C and G were ranked at the 66th percentile, and D and H represented the highest expected difficulty (Fig. 8A). Over the course of two sessions, participants performed the *Yellow Cab* task in all eight environments, encountering one environment from each difficulty rating per session. We counterbalanced the orders in which participants encountered the eight environments. (We found no reliable order effects across the counterbalanced conditions.)

After navigating to each store, participants were instructed to drive away in any direction they wished. After traveling one block from the center of the just-visited store, they were asked to orient a pointer towards the next store in the delivery sequence (Fig. 2C). This pointing task gave a fine-grained snapshot of participants' knowledge before each delivery. The pointing data also helped to validate MAGELLAN's predictions (using parameter values fit to the Experiment 1 data) about which stores were in participants' cognitive maps prior to each delivery. Once a participant pressed a gamepad controller to signal satisfaction with the pointing response, he or she was instructed to navigate to the next store. In addition

to the point-based reward system that we employed in Experiment 1, we used the accuracies of participants' pointing responses to determine their performance-based bonuses. We instructed participants to learn the layouts of the environments, and told them that both their navigation performance and pointing accuracy would determine their monetary bonus.

Validating MAGELLAN

Environment difficulty and learning rate

Like the Experiment 1 participants, the Experiment 2 participants learned to find nearly optimal paths to each store over multiple deliveries (Fig. 3B). As participants gained experience navigating in an environment, they were also able to point more accurately to the remembered locations of stores (Fig. 3C). We next asked whether MAGELLAN's predictions about the relative difficulty of each environment were borne out in the experimental data. A repeated measures analysis of variance (ANOVA) revealed that total excess delivery path length co-varied reliably with MAGELLAN's four difficulty ratings ($F(3, 103) = 34.8, p < 0.001$). This co-variation demonstrates that MAGELLAN-predicted difficulty among the environments was associated with differences in participants' performance in those environments. We next examined the observed total excess path lengths for each environment (Fig. 9A). As can be seen in the panel, participants performed best in environments that had earned an "easy" ranking from MAGELLAN (that is, difficulty level = 1), and their performance dropped as the predicted difficulty rank increased. Post-hoc *t*-tests revealed that participants performed reliably better in environments with difficulty level = 1 than in environments with ranks of 2, 3, or 4 ($t(103) > 8, p$'s < 0.001). Participants also tended to perform better in environments with difficulty level = 2 than in environments with difficulty level = 3 ($t(103) = 1.96, p = 0.05$). The observed total excess path distances for the most difficult environments (difficulty levels = 3 or 4) were not statistically distinguishable from one another ($t(103) = 0.53, p = 0.60$).

We also wondered whether our model accurately predicted the *rates* at which different environments were learned, in addition to participants' overall performance. We fit power law functions to the mean observed and predicted learning curves for each of the eight environments. As a measure of how rapidly the environments were learned, we computed the slopes of these functions in log-log space (where higher absolute values indicate more rapid learning). We found a reliable correlation between the observed and predicted learning rates (Fig. 9B; $r = 0.80, p = 0.02$). Taken together, these outcomes demonstrate that MAGELLAN can make accurate predictions about the ease with which human participants will learn to navigate an environment efficiently. This result is particularly impressive, as the learning rates for each environment were not considered when ranking environments by difficulty. In fact, the estimated learning rates for the eight environments produce a different ordering of the environments than MAGELLAN's difficulty ratings that we used to select the environments from the original set of 500 (correlation between predicted learning rate and expected difficulty: $r = 0.06, p = 0.17$). Nonetheless, MAGELLAN's predictions about the rapidity with which participants would learn to navigate different environments efficiently were reassuringly accurate.

Like expected difficulty, the MAGELLAN-predicted learning rates across the set of 500 randomly generated environments were correlated with the mean distances between successive deliveries ($r = 0.10, p = 0.02$). However, the predicted learning rates were not reliably correlated with the mean distances between stores ($r = 0.06, p = 0.17$). This reflects the fact that the relation between learning rate and environment layout depends on complex interactions between many environment features, the sequence of selected target stores, and moment-by-moment spatial knowledge. MAGELLAN provides a unified framework for taking these interactions into account.

Results of the pointing task

The pointing task in Experiment 2 allowed us to explicitly test MAGELLAN's predictions about which stores were present in participants' cognitive maps prior to each delivery, given their past experience in the environment. If MAGELLAN's cognitive map predictions were accurate, then stores not in the model's cognitive map should be associated with larger pointing errors than stores that were in the model's cognitive map prior to delivery to that store. Indeed, we found this to be the case (Fig. 10). We also found that the pointing errors decreased reliably with the predicted memory strengths associated with the stores (determined by the number of 1-block steps since they were added to memory; circular-linear correlation (Fisher, 1993): $r = -0.19, p < 0.001$). In this way, MAGELLAN's predictions about which landmarks were stored in its cognitive map appeared to accurately mirror the participants' internal cognitive maps, at least to the extent that we could actually probe those maps. This suggests that the MAGELLAN model makes accurate predictions about the evolving states of participants' cognitive maps as they are gaining experience in each environment. Note that purely random guessing would lead to pointing errors of 90° on average (the minimum possible error is 0° and the maximum error is 180°). The fact that participants' pointing errors are far less than 90° even when the corresponding predicted memory strength is 0 (Fig. 10A) suggests that participants are relying on spatial knowledge about other landmarks in the environment to eliminate locations at which they know the target store is *not*; this strategy would be analogous to MAGELLAN's strategy for locating unknown stores by traveling to nearby locations not in the model's cognitive map. We tested whether participants might be relying on this strategy by computing the mean angle to all locations in the model's cognitive map associated with a memory strength of 0 (i.e. all unknown locations) prior to each delivery (Fig. 10C). We found that the pointing errors generated using this strategy were reliably correlated with the pointing errors participants made (circular-circular correlation: $\rho = 0.14, p < 10^{-4}$).

In addition to using pointing data to test MAGELLAN's predictions about which landmarks were stored in participants' cognitive maps, we used the pointing data to test the model's assumption that people use allocentric information to navigate and orient themselves in virtual environments. Following the logic of Kelly, Avraamides, and Loomis (2007) and Mou, Zhao, and McNamara (2007), we reasoned that if people's mental representations of the virtual environments were purely egocentric (i.e., referenced relative to the participants' instantaneous positions and headings), then they should be able to point to landmarks equally well regardless of their absolute heading relative to fixed landmarks in the environment. If, however, people's mental representations also rely on allocentric

coordinates (i.e., relative to global features of the environment that define its principal axes), then they should be able to more easily orient themselves if they point while aligned with these principal axes (i.e., North-South or East-West). Imagine that a participant was not aligned with the environment's principal axes while attempting to point to some landmark. If the participant first mentally rotates his or her heading to align with the environment's principal axes prior to pointing to the landmark (although we did not instruct participants to do so), then the additional cognitive processing required by this operation could introduce additional uncertainty into the pointing and thereby diminish its accuracy. Consistent with this hypothesis, we found that participants pointed to landmarks more accurately when they were oriented within 10° of the environments' principal axes (Fig. 11; Watson-Williams two-sample test (Fisher, 1993): $F = 8.12$, $p = 0.004$).

General Discussion

We introduced MAGELLAN, a high-level allocentric cognitive map-based model of human spatial navigation. The model makes detailed, testable predictions about the ways in which people learn to navigate in unfamiliar environments. We showed that our model accounts for mean performance in Experiment 1 when appropriate adjustments are made to the model's vision and memory parameters. We then used those two parameter values to design environments varying in the expected difficulty participants would have in learning to navigate efficiently. Experiment 2 tested and confirmed MAGELLAN's predictions about the ease with which participants would be able to navigate the environments, and about participants' learning rates across environments. The second experiment also tested MAGELLAN's ability to explain which landmarks were in each participant's cognitive map at various times during the experiment, and we found that the memory strength assigned by MAGELLAN to a given landmark corresponded to the participant's ability to accurately point to it. In this way, MAGELLAN provides a means of designing environments that are particularly easy (or difficult) to learn to navigate in, and the model also provides useful estimates of the evolving states of participants' cognitive maps as they gain experience in an environment.

In the remainder of this section, we first discuss the relation between MAGELLAN and previous computational models of spatial navigation. We then discuss several limitations and potential extensions of the MAGELLAN model. We also provide a brief overview of previous studies of human wayfinding using virtual reality experiments.

Models of spatial navigation

MAGELLAN draws inspiration from several predecessors in the recent navigation literature. These models fall into two general classes: low-level biologically inspired models derived from known neural machinery including place and grid cells (e.g. Burgess & O'Keefe, 1996; Blum & Abbott, 1996; Gerstner & Abbott, 1997; Solstad, Moser, & Einevoll, 2006; McNaughton, Battaglia, Jensen, Moser, & Moser, 2006; Moser et al., 2008), and high-level descriptive models based on egocentric and allocentric spatial encoding strategies that use cognitive maps and path integration (e.g. McNaughton et al., 2006; Han et al., 2012). Whereas biologically inspired models elegantly account for the idea that navigationally relevant cells (Moser et al., 2008; Ekstrom et al., 2003; Jacobs et al., 2013) can support an

allocentric cognitive map, it is not clear how such models would be extended to explain complex navigation behaviors in real or virtual environments. By contrast, while high-level strategy-based models seek to explain complex navigation behaviors, the existing models do not attempt to make specific quantitative predictions about the path a participant will take towards a target given their past experience, or about how spatial information is acquired and stored as a participant navigates in an environment, or about which properties of an environment or spatial task will affect people's abilities to learn to navigate efficiently. MAGELLAN operates at the same high level as strategy-based models, but makes quantitative predictions about the way in which people build up the mental representations of unfamiliar environments and use those representations to navigate efficiently.

A number of recent studies have begun to examine the precise nature of our spatial knowledge concerning an environment, for example by asking whether our cognitive maps are organized in an egocentric (i.e., referenced to the navigator's viewpoint) or allocentric (i.e., referenced by fixed landmarks or orientations in the environment) manner (Klatzky et al., 1990; Wang & Spelke, 2000, 2002; Holmes & Sholl, 2005; Burgess, 2006; Waller & Hodgson, 2006; Avraamides & Kelly, 2010; Byrne & Crawford, 2010; Chen, Byrne, & Crawford, 2011). Our discovery that participants can point to objects in an environment more quickly and accurately when their own viewpoint is aligned with one of the environment's principal axes (e.g. North-South or East-West) supports the notion that a navigator's representation of the environment has at least some allocentric component (Kelly et al., 2007; Mou et al., 2007). Moreover, this result is consistent with other evidence that spatial knowledge is strongly influenced by egocentric variables, including a navigator's orientation and view of an environment, and by the specific paths that the navigator has previously taken (e.g. Andersen, Essick, & Siegel, 1985; Rolls & Xiang, 2006). Our decision to implement an allocentric cognitive map-based model was not intended to suggest that egocentric navigation is less important (or less prevalent) than allocentric navigation. Rather, our goal was to explore the extent to which the process of learning to navigate in a new environment could be explained by a purely allocentric cognitive map representation.

Our work on MAGELLAN also draws inspiration from ideal observer models that have greatly informed our understanding of human perception (for review see Geisler, 2011). Ideal observer models posit that performance is optimal given the knowledge an observer's experience affords them, and that errors or inefficiencies are the consequences of knowledge limitations. The MAGELLAN model uses the ideal observer approach to gain insights into the sources of errors and inefficiencies in spatial cognition and spatial navigation. We found that the ideal navigator outperformed human participants. However, by introducing limitations to the encoding and retention of landmarks, we were able to explain participants' spatial learning curves and the variability in difficulty imposed by an environment's layout and delivery sequence. Our approach is based on the fundamental assumption that variability in the information that a person could possibly have at a given point of learning and in a given environment plays an important role in accounting for variability in behavior.

In a related approach, Stankiewicz, Legge, Mansfield, and Schlicht (2006) constructed an ideal navigator model to study the behaviors of participants who were navigating virtual indoor mazes. Whereas in our visually rich and highly varied towns the main source of

difficulty was a lack of spatial knowledge about the environment, Stankiewicz et al. (2006)'s task involved navigating in visually impoverished environments whose scene information did not provide a unique signal regarding subjects' location and orientation. Their ideal navigator model assumed that observers possess perfect spatial knowledge of the environment (i.e., the scene observed at every position and from every angle) but that difficulty in navigating arose from uncertainty over the observer's location within the environment (rather than a lack of knowledge of landmarks' locations). In this way, whereas MAGELLAN is concerned with acquisition of spatial knowledge, Stankiewicz et al. (2006)'s model is concerned with the problem of spatial updating (i.e., keeping track of one's location within an environment as one experiences varied visual information during navigation). One could easily envision an extension of MAGELLAN that proposes an imperfect process of spatial updating and allows one to determine the degree to which incorrect spatial updating may explain variability in the way people navigate through more visually impoverished environments.

Extending MAGELLAN

Our primary goal in the present study was to identify the principal dimensions of information processing that support experience-based improvements in navigation. Although navigation relies on a large number of highly complex neurocognitive processes, we chose not to construct a complex or highly detailed model of any one process. Rather, MAGELLAN's three modules comprise reduced-form models of how navigators acquire, store, and utilize spatial information. By combining these simple reduced-form models into a single model, we found we were able to provide a useful account of how human navigators build up spatial representations of unfamiliar environments. We next suggest some important potential modifications and extensions to our model aimed at improving its realism and generalizability to a more diverse array of spatial environments.

Vision parameter— V . MAGELLAN's vision parameter is intended to account for the aggregate effects of perceptual and cognitive limitations on acquiring spatial information from the environment. In its current implementation, MAGELLAN assumes that any structure that occupies a sufficiently large fraction of the display will be seen and entered into memory, qualifying the structure as a navigationally relevant landmark (Han et al., 2012). As the model takes no account of a structure's location, a structure far from fixation would be seen and therefore become a landmark so long as it was large enough to exceed V . Because people tend to look in the general direction in which they are walking or driving (Cutting, Readinger, & Wang, 2002; Readinger, Chatziastros, Cunningham, Bühlhoff, & Cutting, 2002), and also because people tend to attend to objects lying in that direction, MAGELLAN's disregard of the relationship between a structure's location and the locus an observer fixates serves to break the customary relationship between location and attention. A more realistic V parameter might explicitly factor in an object's eccentricity by scaling its impact according to a rule parameterized for stimulus eccentricity (e.g. Rovamo & Virsu, 1979). Additionally, although implementing the V parameter as a single thresholded, all-or-none phenomenon serves as a useful simplification, a graded, probabilistic implementation may be more realistic.

Koch and Ullman (1985) noted how local differences in image statistics can guide eye fixations during search of complex scenes, including scenes presented in virtual reality (Pomplun, 2006). In particular, fixations in complex scenes specifically avoid areas that are uninformative, where informativeness is defined in terms of scene statistics and task-relevance (Kayser, Nielsen, & Logothetis, 2006). It is likely that future extensions of MAGELLAN would gain in predictive power by incorporating parameters that reflect image-based differences in salience (Peters, Iyer, Itti, & Koch, 2005), as well as parameters that represent task-dependent variations in looking behaviors (Hayhoe, Bensinger, & Ballard, 1998; Pelz & Canosa, 2001).

Finally, MAGELLAN has no representation of visual similarity, a variable that has considerable influence in visual processing. Rather, the model assumes that any landmark that is seen and entered into spatial memory will be distinct from any other landmark that might be seen and entered into memory. As a result, MAGELLAN can retrieve the location of any stored landmark without error, never confusing one previously seen landmark with another. In the two experiments reported here, this assumption of completely distinctive landmarks seems justified. After all, participants were able to give a verbal label to each store that was seen, and the store names in any single environment were designed to be highly distinctive. Additionally, prior to testing, participants were familiarized with the name and appearance of every store that might be encountered. It seems likely, though, that in real-world wayfinding, people are influenced by the visual distinctiveness or differentiation of landmarks. City planners and researchers have long known that environments whose elements are differentiated (e.g. with regard to architectural style or color) are easier to wayfind in than more homogeneous environments. Montello (2005) gives an excellent review of the characteristics of physical environments that influence orientation during navigation. The gain in predictive power that might accrue from any of these plausible ways to define a more nuanced, non-scalar version of MAGELLAN's V parameter would need to be evaluated relative to additional degrees of freedom they would require.

Memory parameter— M . For the sake of simplicity, M comprised a surrogate for what are likely to be several different operations. One potential elaboration of M would allow variability in the number of steps before any particular memory is no longer viable. It seems reasonable that some landmarks are “forgotten” more quickly than others, perhaps as a function of depth of processing and/or as a function of the landmarks' distinctiveness (Montello, 2005). Further, we should note that MAGELLAN is not committed to an all-or-none forgetting process. Rather, we could easily envision spatial location information being lost gradually as new, interfering information is learned. Moreover, information that has become inaccessible might later be retrieved given a salient contextual or associative cue (McGeoch, 1932; Tulving & Pearlstone, 1966; Howard & Kahana, 2002).

Ideal navigator and model design—Our treatment of an ideal navigator made an assumption about the utility function that the navigator would attempt to minimize. Specifically, we assumed that navigators would seek to minimize distance traveled, an assumption that is consistent with the structure of our task—with participants' bonus pay

contingent upon excess path length. However, one can easily envision equally realistic driving tasks in which participants attempted to minimize travel time rather than distance. Alternatively, participants might wish to minimize travel along particular kinds of roads, such as toll roads, or they may plan routes that afford scenic views, that are expected to have little traffic, or that are more familiar.

Virtual reality studies of wayfinding

Our study concentrated on navigation in virtual environments. Therefore, it is reasonable to ask whether our result can be extrapolated to environments other than virtual ones. Over the last two decades, virtual reality (VR) technology has made possible a wide range of applications, including video games, training for the military and first responders, desensitization therapies for phobias, and regimens for rehabilitation after brain injury (Zyda, 2005; Satava, 1995; Botella et al., 1998; Grealy, Johnson, & Rushton, 1999). Technological advances make it possible to generate customized, interactive VR environments within which interactions can be manipulated, measured, and analyzed. Numerous previous studies have thus used VR tasks to study both the cognitive and neural basis of human spatial navigation (Newman et al., 2007; Korolev, Jacobs, Mollison, & Kahana, 2005; Jacobs et al., 2007; Manning et al., 2009; Jacobs et al., 2010, 2010, 2013). Even very compelling VR is not likely to be confused with the real world, but so long as a VR environment can claim sufficient realism, knowledge gained from a study set in VR can be extrapolated to the real world (Whitmer, Bailey, Knerr, & Parsons, 1996). Furthermore, previous studies by Ekstrom et al. (2003), Jacobs et al. (2010), and Jacobs et al. (2013) showed that similar VR towns to those used in Experiment 1 of this paper were sufficiently realistic to activate navigationally useful place, view, goal, and grid cells in the brains of humans who were navigating through those towns, and similar VR environments have also been shown to exhibit grid-cell like responses in human fMRI data (Doeller, Barry, & Burgess, 2010).

Of course it is difficult to know *a priori* when enough realism is enough, or exactly what attributes a VR environment must afford to navigators in a laboratory study. However, experimental tests of hypotheses can provide useful clues. For instance, we know that navigation performance degrades when the illusion of natural movement is removed, for example by reducing the rate of optic flow a participant experiences in a VR environment (Kirschen, Kahana, Sekuler, & Burack, 2000). We also know that distinctive building textures and sizes are important features of a VR environment. In fact, when an environment's texture is overly homogeneous, participants frequently become lost, and tend to fall back upon simple list learning strategies, such as "turn left, then right turn, etc.," rather than learning the environment's spatial layout (Kirschen et al., 2000). Counterintuitively, large-field displays and 3-dimensional projections do not seem to enhance people's ability to learn to navigate in virtual environments (Dahmani, Ledoux, Boyer, & Bohbot, 2011).

Using VR, researchers can build multiple environments from a common set of generative rules so that performance can be compared across those environments, as we have done. By using explicit rules to govern the layouts of multiple environments and the placements of

landmarks within those environments, it is possible to identify the principles that govern navigation and learning, independent of the idiosyncratic navigational challenges that are posed by any single environment (Newman et al., 2007). In this way, results from studies of wayfinding in properly designed VR environments are arguably more generalizable than results from any single real world environment.

Concluding remarks

Although our two-parameter model gave a good account of performance in both experiments, what has been achieved is clearly only a step towards a more complete account of spatial learning and navigation. For example, a comprehensive model of spatial learning and spatial navigation should account for learning of route-based information (landmark-to-landmark associations) alongside the learning of map-based information (spatial location to landmark associations) that we emphasize in the MAGELLAN model. The importance of route-based coding is clear from the considerable evidence for orientation dependence in human spatial cognition (e.g. Shelton & McNamara, 2001; Mou, Zhang, & McNamara, 2004). Both Benhamou, Bovet, and Poucet (1995) and Schölkopf and Mallot (1995) proposed route-based models of spatial information processing based on landmark-to-landmark rather than position-to-landmark associations. These ideas have a clear parallel in the sequence learning literature, in which Ladd and Woodworth (1911) first proposed that both position-to-item associations *and* chained, item-to-item associations are important in serial learning (see also, Young, 1968; Lewandowsky & Murdock, 1989; Burgess & Hitch, 1999, 2005; Brown, Preece, & Hulme, 2000).

MAGELLAN makes three main contributions to the literature on spatial navigation. First, MAGELLAN's ability to explain the rate of experience-based learning in both experiments suggests that human navigation may be largely supported by allocentric cognitive map representations (e.g. perhaps supported by populations of place and grid cells). Second, MAGELLAN provides a way to estimate the evolution of participants' cognitive maps over time, as a previously unfamiliar environment is being learned. Third, MAGELLAN introduces a quantitative benchmark against which spatial navigation behavior can be evaluated. To generate such a benchmark we adjusted MAGELLAN's V and M parameters to account for mean performance in Experiment 1, using the same parameters for all participants. However, fitting these parameters to individual participants' performance could generate important insights into the character of individual differences in wayfinding ability (e.g. see Kozlowski & Bryant, 1977). Once such fits were obtained, one could predict how individual participants' cognitive maps would vary over the course of the experiment. This would be particularly valuable as it would support comparisons among the distributions of best-fitting parameters obtained for different populations (e.g. younger vs. older adults, healthy vs. impaired individuals, etc.).

MAGELLAN formalizes cognitive map-based theories of spatial cognition by estimating how the states of participants' cognitive maps evolve as they gain experience navigating in a new environment. By showing that the model accurately estimates how people learn to navigate more efficiently with experience, how difficult it will be to learn to navigate efficiently in different environments, and which landmarks will be viable in participants'

cognitive maps, we show that MAGELLAN provides a useful tool for understanding and assessing participants' abilities to learn to navigate in unfamiliar environments.

Acknowledgments

Supported by NIH grants MH61975, MH55687, MH068404, and F31MH088118, and by NSF grants SBE-0354378 and BCS-0527689. We are grateful for the assistance of and/or useful discussions with Gaurav Bharaj, Arthur Hicks, Josh Jacobs, Igor Korolev, Talia Manning, Matt Mollison, Ben Okaty, Andrew Spellman, and Christoph Weidemann. The content is solely the responsibility of the authors and does not necessarily represent the official views of our supporting organizations. The experimental paradigms, data, analysis code, and a MATLAB implementation of the MAGELLAN model may be downloaded from <http://memory.psych.upenn.edu>.

References

- Andersen R, Essick G, Siegel R. Encoding of spatial location by posterior parietal neurons. *Science*. 1985; 230(4724):456–458. [PubMed: 4048942]
- Avraamides MN, Kelly JW. Multiple systems of spatial memory: evidence from described scenes. *Journal of Experimental Psychology: Learning Memory and Cognition*. 2010; 36(3):635–645.
- Benhamou S, Bovet P, Poucet B. A model for place navigation in mammals. *Journal of Theoretical Biology*. 1995; 173:163–178.
- Bennett PJ, Pratt J. The spatial distribution of inhibition of return. *Psychological Science*. 2001; 12(1): 76–80. [PubMed: 11294232]
- Blum KI, Abbott LF. A model of spatial map formation in the hippocampus of the rat. *Neural Computation*. 1996; 8:85–93. [PubMed: 8564805]
- Botella C, Banos RM, Perpina C, Villa H, Alcaniz M, Rey A. Virtual reality treatment of claustrophobia: a case report. *Behavior research and therapy*. 1998; 36:239–246.
- Brown GDA, Preece T, Hulme C. Oscillator-based memory for serial order. *Psychological Review*. 2000; 107(1):127–181. [PubMed: 10687405]
- Burgess N. Spatial memory: How egocentric and allocentric combine. *Trends in Cognitive Science*. 2006; 10(12):551–557.
- Burgess N, Hitch G. Computational models of working memory: putting long-term memory into context. *tics*. 2005; 9:535–41.
- Burgess N, Hitch GJ. Memory for serial order: A network model of the phonological loop and its timing. *Psychological Review*. 1999; 106(3):551–581.
- Burgess N, O'Keefe J. Neuronal computations underlying the firing of place cells and their role in navigation. *Hippocampus*. 1996; 7:749–762. [PubMed: 9034860]
- Byrne PA, Crawford JD. Cue reliability and a landmark stability heuristic determine relative weighting between egocentric and allocentric visual information in memory-guided reach. *Journal of Neurophysiology*. 2010; 103:3054–3069.
- Caplan JB, Madsen JR, Schulze-Bonhage A, Aschenbrenner-Scheibe R, Newman EL, Kahana MJ. Human theta oscillations related to sensorimotor integration and spatial learning. *Journal of Neuroscience*. 2003; 23:4726–4736. [PubMed: 12805312]
- Chen Y, Byrne P, Crawford JD. Time course of allocentric decay, egocentric decay, and allocentric-to-egocentric conversion in memory-guided reach. *Neuropsychologia*. 2011; 49:49–60. [PubMed: 21056048]
- Cutting JE, Readinger WO, Wang RF. Walking, looking to the side, and taking curved paths. *Perception & Psychophysics*. 2002 Apr; 64(3):415–425. [PubMed: 12049282]
- Dahmani L, Ledoux AA, Boyer P, Bohbot V. Wayfinding: the effects of large displays and 3-D perception. *Behavior Research Methods*. 2011; 43(1):1–8. [PubMed: 21287104]
- Doeller CF, Barry C, Burgess N. Evidence for grid cells in a human memory network. *Nature*. 2010; 463(7281):657–661. [PubMed: 20090680]
- Ekstrom AD, Caplan J, Ho E, Shattuck K, Fried I, Kahana M. Human hippocampal theta activity during virtual navigation. *Hippocampus*. 2005; 15:881–889. [PubMed: 16114040]

- Ekstrom AD, Kahana MJ, Caplan JB, Fields TA, Isham EA, Newman EL, et al. Cellular networks underlying human spatial navigation. *Nature*. 2003; 425:184–187. [PubMed: 12968182]
- Fisher, NI. *Statistical analysis of circular data*. Cambridge, England: Cambridge University Press; 1993.
- Fried I, Wilson C, Maidment N, Engel JJ, Behnke E, Fields T, et al. Cerebral microdialysis combined with single-neuron and electroencephalographic recording in neurosurgical patients. *Journal of Neurosurgery*. 1999; 91:697–705. [PubMed: 10507396]
- Geisler WS. Sequential ideal-observer analysis of visual discriminations. *Psychological Review*. 1989; 96(2):267– 314. [PubMed: 2652171]
- Geisler WS. Contributions of ideal observer theory to vision research. *Vision Research*. 2011; 51:771–781. [PubMed: 20920517]
- Geller AS, Schleifer IK, Sederberg PB, Jacobs J, Kahana MJ. PyEPL: A cross-platform experiment-programming library. *Behavior Research Methods*. 2007; 39(4):950–958. [PubMed: 18183912]
- Gerstner W, Abbott LF. Learning navigational maps through potentiation and modulation of hippocampal place cells. *Journal of Computational Neuroscience*. 1997; 4:79– 94. [PubMed: 9046453]
- Grealy MA, Johnson DA, Rushton SK. Improving cognitive function after brain injury: the use of exercise and virtual reality. *Archives of Physical Medicine and Rehabilitation*. 1999; 80(6):661–667. [PubMed: 10378492]
- Han X, Byrne P, Kahana MJ, Becker S. When do objects become landmarks? A VR study of the effect of task relevance on spatial memory. *PLoS One*. 2012; 7(5):e35940. [PubMed: 22586455]
- Hayhoe MM, Bensinger DG, Ballard DH. Task constraints in visual working memory. *Vision Research*. 1998; 38(1):125– 137. [PubMed: 9474383]
- Holmes MC, Sholl MJ. Allocentric coding of object-to-object relations in overlearned and novel environments. *Journal of Experimental Psychology: Learning Memory and Cognition*. 2005; 31(5):1069– 1087.
- Howard MW, Kahana MJ. A distributed representation of temporal context. *Journal of Mathematical Psychology*. 2002; 46:269–299.
- Huth, JE. *The lost art of finding our way*. The President and Fellows of Harvard College; 2013.
- Jacobs J, Kahana MJ, Ekstrom AD, Fried I. Brain oscillations control timing of single-neuron activity in humans. *Journal of Neuroscience*. 2007; 27(14):3839–3844. [PubMed: 17409248]
- Jacobs J, Kahana MJ, Ekstrom AD, Mollison MV, Fried I. A sense of direction in human entorhinal cortex. *Proceedings of the National Academy of Sciences, USA*. 2010; 107(14):6487–6482.
- Jacobs J, Weidemann CT, Miller JF, Solway A, Burke JF, Wei X, et al. Direct recordings of grid-like neuronal activity in human spatial navigation. *Nature Neuroscience*. 2013; 16:1188–1190.
- Kayser C, Nielsen KJ, Logothetis NK. Fixations in natural scenes: interaction of image structure and image content. *Vision Research*. 2006; 46(16):2535– 2545. [PubMed: 16545420]
- Kelly JW, Avraamides MN, Loomis JM. Sensorimotor alignment effects in the learning environment and in novel environments. *Journal of Experimental Psychology*. 2007; 33(6):1092– 1107. [PubMed: 17983315]
- Kirschen MP, Kahana MJ, Sekuler R, Burack B. Optic flow helps humans learn to navigate through synthetic environments. *Perception*. 2000; 29:801–818. [PubMed: 11064802]
- Klatzky RL, Loomis JM, Golledge RG, Cicinelli JG, Doherty S, Pellegrino JW. Acquisition of route and survey knowledge in the absence of vision. *Journal of Motor Behavior*. 1990; 22(1):19– 43. [PubMed: 15111279]
- Klein RM. Inhibition of return. *Trends in Cognitive Science*. 2000; 4(4):138– 147.
- Koch C, Ullman S. Shifts in selective visual attention: Towards the underlying neural circuitry. *Human Neurobiology*. 1985; 4(4):219– 227. [PubMed: 3836989]
- Korolev IO, Jacobs J, Mollison MV, Kahana MJ. Human oscillatory activity during virtual navigation: a comparison between scalp and intracranial recordings. *Society for Neuroscience Abstracts*, Program No. 65.16. 2005
- Kozlowski L, Bryant K. Sense of direction, spatial orientation, and cognitive maps. *Journal of Experimental Psychology: Human Perception and Performance*. 1977; 3(4):590–598.

- Ladd, GT.; Woodworth, RS. Elements of physiological psychology: A treatise of the activities and nature of the mind from the physical and experimental point of view. New York, NY: Charles Scribner's Sons; 1911.
- Lewandowsky S, Murdock BB. Memory for serial order. *Psychological Review*. 1989; 96:25–57.
- Lyon DR, Gunzelmann G, Gluck KA. A computational model of spatial visualization capacity. *Cognitive Psychology*. 2008 Sep; 57(2):122–152. [PubMed: 18314098]
- Manning JR, Jacobs J, Fried I, Kahana MJ. Broadband shifts in LFP power spectra are correlated with single-neuron spiking in humans. *Journal of Neuroscience*. 2009; 29(43):13613– 13620. [PubMed: 19864573]
- McGeoch JA. Forgetting and the law of disuse. *Psychological Review*. 1932; 39:352–70.
- McNaughton BL, Battaglia FP, Jensen O, Moser EI, Moser MB. Path integration and the neural basis of the 'cognitive map'. *Nature Reviews Neuroscience*. 2006; 7:663–678. Retrieved from <http://dx.doi.org/10.1038/nrn1932>.
- Montello, D. The cambridge handbook of visuospatial thinking. Shah, P.; Miyake, A., editors. Cambridge, U. K: Cambridge University Press; 2005. (chap. Navigation)
- Moser E, Kropff E, Moser M. Place cells, grid cells, and the brain's spatial representation system. *Annu Rev Neurosci*. 2008; 31:69–89. [PubMed: 18284371]
- Mou W, Zhang K, McNamara TP. Frames of reference in spatial memories acquired from language. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2004; 30(1):171–180.
- Mou W, Zhao M, McNamara T. Layout Geometry in the Selection of Intrinsic Frames of Reference from Multiple Viewpoints. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2007; 33(1):145–54.
- Müller M, Wehner R. The hidden spiral: systematic search and path integration in desert ants, *cataglyphis fortis*. *J Comp Physiol A*. 1994; 175:525– 530.
- Newman EL, Caplan JB, Kirschen MP, Korolev IO, Sekuler R, Kahana MJ. Learning your way around town: How virtual taxicab drivers learn to use both layout and landmark information. *Cognition*. 2007; 104(2):231–253. [PubMed: 16879816]
- O'Keefe J, Dostrovsky J. The hippocampus as a spatial map: Preliminary evidence from unit activity in the freely-moving rat. *Brain Research*. 1971; 34:171–175. [PubMed: 5124915]
- Pelz JB, Canosa R. Oculomotor behavior and perceptual strategies in complex tasks. *Vision Research*. 2001; 41(25):3587– 3596. [PubMed: 11718797]
- Peters RJ, Iyer A, Itti L, Koch C. Components of bottom-up gaze allocation in natural images. *Vision Research*. 2005; 45(18):2397– 2416. [PubMed: 15935435]
- Pomplun M. Saccadic selectivity in complex visual search displays. *Vision Research*. 2006; 46(12): 1886– 1900. [PubMed: 16445960]
- Readinger WO, Chatziastros A, Cunningham DW, Bühlhoff HH, Cutting JE. Gaze eccentricity effects on road position and steering. *Journal of Experimental Psychology: Applied*. 2002 Dec; 8(4):247– 258. [PubMed: 12570099]
- Robles-De-La-Torre G, Sekuler R. Numerically estimating internal models of dynamic virtual objects. *ACM Transactions in Applied Perception*. 2004; 1(2):102– 111.
- Rolls E, Xiang JZ. Spatial view cells in the primate hippocampus and memory recall. *Rev Neuroscience*. 2006; 17:175–200.
- Rovamo J, Virsu V. An estimation and application of the human cortical magnification factor. *Experimental Brain Research*. 1979; 37(3):495–510. [PubMed: 520439]
- Satava RM. Virtual reality and telepresence for military medicine. *Comput Biol Med*. 1995; 25(2): 229– 236. [PubMed: 7554840]
- Schölkopf B, Mallot H. View-based cognitive mapping and path planning. *Adaptive Behavior*. 1995; 3:311–348.
- Shelton AL, McNamara TP. Systems of spatial reference in human memory. *Cognitive Psychology*. 2001; 43:274–310. [PubMed: 11741344]
- Solstad T, Moser EI, Einevoll GT. From grid cells to place cells: A mathematical model. *Hippocampus*. 2006; 16:1026–1031. [PubMed: 17094145]

- Solway A, Miller JF, Kahana MJ. PandaEPL: A library for programming spatial navigation experiments. *Behavior Research Methods*. 2013;1–20. [PubMed: 23055156]
- Stankiewicz BJ, Legge GE, Mansfield JS, Schlicht EJ. Lost in virtual space: studies in human and ideal spatial navigation. *Journal of Experimental Psychology: Human Perception and Performance*. 2006; 32(3):688–704.
- Tolman EC. Cognitive maps in rats and men. *Psychology Review*. 1948; 55:189–208.
- Tulving E, Pearlstone Z. Availability versus accessibility of information in memory for words. *Journal of Verbal Learning and Verbal Behavior*. 1966; 5:381–391.
- Waller D, Hodgson E. Transient and enduring spatial representations under disorientation and self-rotation. *Journal Experimental Psychology: Learning, Memory and Cognition*. 2006; 32(4):867–882.
- Wang RF, Spelke ES. Updating egocentric representations in human navigation. *Cognition*. 2000; 77:215–250. [PubMed: 11018510]
- Wang RF, Spelke ES. Human spatial representation: insights from animals. *Trends Cogn Sci*. 2002; 6(9):376–382. [PubMed: 12200179]
- Watrous AJ, Fried I, Ekstrom AD. Behavioral correlates of human hippocampal delta and theta oscillations during navigation. *Journal of Neurophysiology*. 2011; 105(4):1747–1755. [PubMed: 21289136]
- Whitmer BG, Bailey JH, Knerr BW, Parsons KC. Virtual spaces and real-world places: Transfer of route knowledge. *International Journal of Human-Computer Studies*. 1996; 45:413–428.
- Young, RK. Serial learning. In: Dixon, TR.; Horton, DL., editors. *Verbal behavior and general behavior theory*. Englewood Cliffs, N. J: Prentice-Hall; 1968. p. 122-148.
- Zyda M. From visual simulation to virtual reality games. *Computer*. 2005; 38(9):25–32.




Figure 1. The MAGELLAN model

The *vision* module is the model's means for acquiring new information about the visual environment. Structures that occupy at least $V \times 100$ percent of the screen are added to the model's cognitive map. The *cognitive map* module stores acquired spatial information. In the diagram, potential targets are represented in red, and other landmarks are represented in black. Darker colored squares correspond to "stronger memories" of the landmark at the corresponding location, while white squares represent locations whose structures are not encoded in the spatial memory. The navigator's current location and heading are indicated by the yellow teardrop, and the location of the passenger is indicated by the blue star. Memories stored in the cognitive map remain viable for M steps. A *route generation* module computes an efficient route between locations stored in the cognitive map, taking account of the environment's city-block structure. When a target destination is not contained in the spatial memory, the navigator's current location is compared to "blank" blocks on the cognitive map (shown in white). MAGELLAN's route generation module then produces an efficient path to the nearest unknown block. As the path is navigated, the vision module feeds new data to the cognitive map, which is used to update the route in turn. In this way, each module feeds information to the next.



Figure 2. Screen captures from virtual navigation experiments

A. A participant's view as they search for a passenger to pick up in Experiment 1. The rectangles (not visible to the participant during the experiment) illustrate 1%, 8%, and 16% of the screen. **B.** A participant's view as they search for the target store, Pizzeria, in Experiment 2. **C.** A participant's view as they orient an onscreen pointer towards Garden Store in Experiment 2.




Figure 3. Spatial learning curves

A. Mean excess path distance as a function of the number of passengers that have been delivered (Experiment 1). Error bars denote \pm SEM. **B.** Spatial learning curve for Experiment 2. Negative excess path length indicates that the ideal path distance (equal to the city block distance between the passenger and goal) is greater than the delivery path distance (the summed Euclidean distance between adjacent points along the participant's delivery path). **C.** Pointing error as a function of destination number for Experiment 2 participants. Error bars denote \pm circular SEM (Fisher, 1993).




Figure 4. Ideal Navigator predictions for Experiment 1
The curves are in the same format as those in Figure 3A and B.




Figure 5. MAGELLAN's best fits to spatial learning curves from Experiment 1
 Curves are in the same format as those shown in Figure 3A,B. Here we used $(V, M) = (0.08, 32)$, $\text{RMSD} = 0.652$ blocks. V determines a threshold fraction of the screen which a landmark must occupy in order to be added to the spatial memory. M defines the number of steps taken before a new memory in the model's cognitive map becomes unusable.



Figure 6. Examples of buildings and stores that were used to construct Experiment 2 structures
A. Building textures were automatically generated from a set of four brick textures, five window styles, four door styles, and three door placement options, as shown. **B.** After assembling the full texture for a building, the texture image was cropped as needed and hue-shifted while maintaining the covariance structure of pixel intensities. **C.** Eight store images are shown. A total of 44 unique store images were used in the experiment (5 per environment, with an additional 4 stores used in an initial practice task intended to familiarize participants with the use of the gamepad controller). Figure 2B shows one fully rendered Experiment 2 town.




Figure 7. Store spread in Experiment 1 and Experiment 2 environments

A. Smoothed distribution of mean distances between stores in Experiment 1 (dark gray) and Experiment 2 (light gray) environments. **B.** Smoothed distribution of mean distances between successive deliveries in Experiment 1 and 2 (same color scheme as Panel A).




Figure 8. Expected difficulties of randomly generated environments

We randomly generated 500 environments and used MAGELLAN to compute each environment's predicted difficulty, defined as the sum of the average simulated excess path lengths across all of the deliveries made for that environment. From the distribution, eight environments, A–H, were selected for use in Experiment 2. **A.** The smoothed distribution of expected difficulties predicted for all 500 environments rated by MAGELLAN. The difficulties predicted for environments A–H are denoted by vertical lines. **B–D.** Environmental factors that correlate with predicted difficulty. Each dot corresponds to a single randomly generated environment. The x -coordinates of each dot denote that environment's predicted difficulty. Pearson's correlations are reported in each panel. **B.** The y -coordinates of the data points reflect the average Euclidean distances between the centers of the stores in the environments; predicted difficulty increases as the stores become more spread out. **C.** The y -coordinates of the data points reflect the average Euclidean distance between stores navigated to in succession. Predicted difficulty increases as the distance between successive goals increases. **D.** The y -coordinates of the data points reflect the average inter-store distance (Panel C) divided by the average store spread (panel B) of each environment. Predicted difficulty increases with inter-store distance, even after taking store spread into account. This indicates that the correlation between predicted difficulty and inter-store distance (shown in Panel C) does not simply reflect the store spread (shown in Panel B).




Figure 9. Predicted vs. observed environment difficulty

A. Observed mean total excess path length for environments from each of the four difficulty rankings (see Fig. 8). Error bars denote \pm SEM. **B.** Observed and predicted learning rates for each of the 8 environments. Learning rates were estimated by fitting power law functions to the mean observed (or predicted) learning curves for each environment. The indicated learning rates are the fitted slopes of these functions in log-log space. Steeper slopes (higher absolute values) indicate more rapid learning.




Figure 10. Pointing error predictions

A. This panel indicates the mean observed pointing error at the start of deliveries for which the goal store was (Strength > 0) or was not (Strength = 0) in MAGELLAN's cognitive map.

B. Here we subdivided trials where the goal was in the cognitive map (Strength > 0) into 10 equally sized bins according to the predicted strength of the memory for the goal in MAGELLAN's cognitive map. We then computed the mean observed pointing error for the trials contained in each bin. The reported circular-linear correlation for the 10 bins is shown in the panel; the text reports the un-binned correlation. Error bars in Panels A and B denote \pm circular SEM (Fisher, 1993). **C.** Estimated and observed pointing errors. We estimated the pointing errors using the state of MAGELLAN's cognitive map prior to each delivery to eliminate the locations of known buildings (Strength > 0); see text for details. Each data point reflects 1% of the observations. The reported circular-circular correlation (Fisher, 1993) reflects the full distributions of estimated and observed pointing errors.




Figure 11. Pointing error as a function of navigators' heading

The bar labelled "Align" represents the mean pointing error (\pm circular SEM) made when participants were oriented within 10° of the environments' principal, North-South or East-West, axes; the bar labelled "Misalign" shows mean pointing error when participants were not oriented near one of the principal axes. Participants tended to point more accurately while aligned with the environments' principal axes, indicating that they rely, at least partially, on allocentric cues.