Where am I? The practical organisation of finding one's location in a geography computer game

Andrei Korbut

Independent scholar

Abstract

The paper provides a descriptive analysis of the practical organisation of playing Geoguessr, a geography computer game in which the player has to find their location on Google Maps using only Google Street View. It is argued that in order to do this, the player assembles an occasioned corpus of geographical knowledge. It is shown that by constantly foregrounding and backgrounding details of the unfolding scenes in search for clues to "where I am", the player anchors common-sense geographical categories in the sequences of enacted actions. Finally, it is proposed that the activity of playing Geoguessr, while not constituting the discovery practice *per se*, shares some commonalities with other discovery practices, such as scientific discovery or detective work.

INTRODUCTION: FINDING OUT "WHERE I AM" FOR FUN

The question "Where am I?" can be a serious one. Most people have had the experience of getting lost and needing to determine their location. Fortunately, it's rarely a matter of life and death, as when you are lost in the desert or at sea. Usually, we wonder where we are in much more mundane circumstances, such as finding a recommended coffee shop or walking through an unfamiliar city. And there are plenty of tools to assist us in these: from printed maps in tourist guidebooks (Laurier and Brown 2008) to mobile phones (Laurier 2001) and navigational software (Laurier, Brown, and McGregor 2016).

Repeated exposure to the problem of finding one's location can lead to the development of embodied systems of geographical knowledge, such as those described by David Lewis in the case of Pacific Ocean navigation (Lewis 1994) and Chukchi wayfinding in tundra and frozen sea (Lewis and George 1991). People learn to orient by stars, trees, animals, winds, clouds, waves, currents, the curvatures of the landscape, and many other features of the environment. In the urban settings, this orientation is typically mediated by various forms of text and signage, as well as landmarks and vehicles. Today, mobile electronic devices and software provide an additional layer of geographic information, but also create new problems re-

lated to the need of finding correspondence between this layer and the real circumstances in which such information is used.

However important is the question of finding your location in serious mundane or professional contexts, it can also be asked for fun, as part of a leisure activity. There are forms of sport, such as orienteering, where asking and answering "where am I?" is a part and parcel of doing sporting (Pehkonen, Smith, and Smith, 2022). Another such leisure activity is computer games. Some of them offer players whole new fantasy worlds to explore. Some engage players in chasing each other in the real world using specialised software (Benford et al. 2006). And some offer the opportunity to explore our own familiar world using existing navigational systems (Smith et al. 2020).

This paper deals with the one such game, a geography computer game called Geoguessr, developed by Anton Wallén. Geoguessr encourages players to explore the "real world" with the help of the game interface and the popular navigational software Google Maps. The idea is quite simple: you are randomly placed in a location on planet Earth and you have to figure out where you are using only Google Street View and the default type of Google Maps with the search function disabled. Here is an example of the initial setting in which the player might find themselves.

Figure 1

The street-level image of an unfamiliar—or, if you are extremely lucky, familiar to you—place is accompanied by the game interface: a compass, zoom buttons, a "Return to Start" button (small round button with a flag in the bottom left corner of the screen), and the option to set the interim point of return. There is also an enlargeable map in the bottom right corner and an information tableau in the top right corner. The controls will be familiar to anyone who used Google Street View: you can move around by clicking on the destination spot, by clicking on white arrows, or by pressing arrow keys on your keyboard; you can zoom in/out with the help of a mouse or keyboard; and you can turn around by dragging the image. Once you are ready to guess where you are, you put a pin on the map and, depending on how close is your guess to the initial position, you get a corresponding number

of points, with the highest score being 5000 for a sufficiently close guess (the distance to the initial position, which is rewarded with 5000 points, is not constant, but depends on the location).

Unlike many other ways of finding where you are for fun, in Geoguessr the player does not have to go outside: all the navigation is done via the computer in the virtual 3D environment built from the images of real places. Instead of figuring out where you are by constantly relating what you see around you with the map, the player stays in front of the computer and tries to determine the geographical location of the initial drop point using the digital images displayed on the screen. However, since Geoguessr's virtual environment imitates the real world and, as we will see later, movement in this environment resembles the movement in the real world, it can be said that the player's navigation in the game is comparable to navigation in the real world.

In this paper I will look at how players of Geoguessr find out where they are using the basic setup of the game. There are several options and game types available to players (the player can set the timer and try to find their location in, say, 2 minutes; they can play in multiplayer mode; etc.), but I will only deal with the simplest mechanic: the player plays five rounds alone with no time limit. I will assume that, doing this, the player does not use any additional tools such as navigation apps, search engines, or road atlases. While no one can stop the player from obtaining information from these sources, this usually ruins the fun.

In what follows, I will first introduce the general framework for my analysis, followed by a descriptive analysis of a complete episode of the player finding where they are. The player in this episode will be me. For this paper, I've selected a particular case where I was able to find where I was with a 3-meter error in 3 minutes 45 seconds. This is my personal record so far in terms of the time taken to find my exact location, but of course it is far behind what skilled players can achieve.[1] In this respect, the game that I am going to analyse reflects the skill level of the novice player. But I will not focus on the development of skills within the game. Instead, I will use this case as a perspicuous setting to highlight key elements of navigational practices in the game.

After the descriptive analysis, I will discuss whether it makes sense to compare searching practices in Geoguessr with discovery practices in science and beyond.

FRAMEWORK: "WHERE AM I?" AS A PRACTICAL ACHIEVEMENT

Before moving on to an analysis of the actual practice of playing Geoguessr, I need to lay out some concepts and distinctions that will define the focus of the proposed analysis and provide the general framework for the discussion.

First of all, although I said earlier that in Geoguessr the player's practice consists of finding where they are, this statement is not without problems and requires a comment. To be more precise, the player finds "where it is" and not "where I am". Contrary to the person who wonders where they are, say, in the middle of a busy street, the player in Geoguessr

1 A good player can make 5k score in 5 seconds! See: https://www.youtube.com/watch?v=ZWq3G7Vdu4c.

knows where they are: they are in front of the computer, be it a desktop computer, a laptop, or a mobile device. They need to find out where is the virtual spot where the game has dropped them. However, the distinction between "where it is" and "where I am" is not clear-cut here. The virtual 3D environment of the game and the ways it provides for the movement in this environment give the impression that the player is moving down the road among urban or rural objects. The game is played in the first person, and the player can do things that resemble what real persons could do in the real setting, such as "looking around". Thus, practically speaking, we can use the questions "Where is it?" and "Where am I?" as interchangeable. But we must remember that "I" here is not an actual body, just as "it" is not an actual place. Both are virtual entities that only make sense in the sequences of actions inside the computer game.

It is also important to unpack the "where" in both questions. There are two components of this "where". On the one hand, in terms of what the player is striving for, the "where" is the spot on the map where the player puts their pin. But on the other hand, this spot is always accompanied by a thematical "positional index" (Gurwitsch 2010, 351), i.e., it occupies a place in what Schegloff calls "common sense geography" or "everyman's geography" (Schegloff 1972, 85), meaning our ordinary structuring of the categories of places. This common-sense geography has a particular organisation that can be depicted as follows.

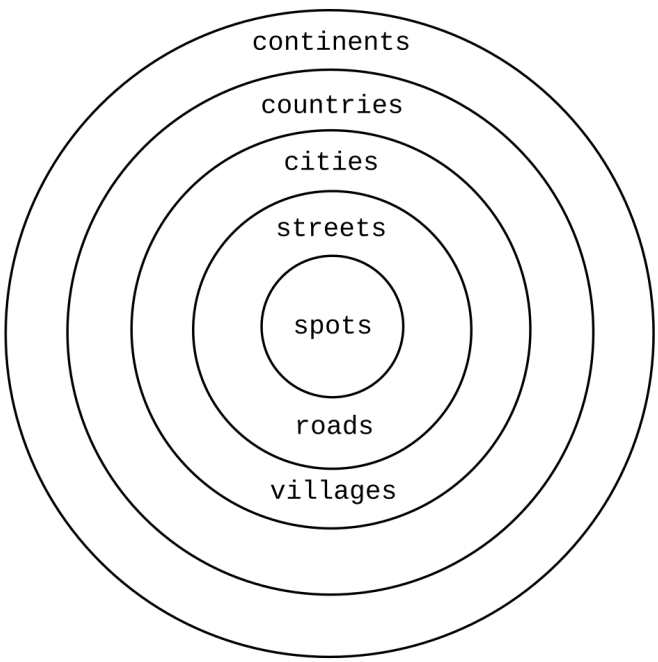

The categories that make up everyman's geography are organised in concentric circles, with one set nested within the other according to their relative scale. Some categories belong to the same circle, although they may be of different sizes (e.g., towns and villages), and there are universal and more idiosyncratic ways of compartmentalizing the categories, such as di-

viding each into western/eastern/northern/southern parts, or distinguishing large and small exemplars of them (e.g., large and small countries).

The ability to zoom in and out in Google Maps (or any other cartographic software) prioritises the concentric organisation of the categories, but we should not let this blind us to the situated work that is required to organise them in this way. Geographical categories are situated: ". . . occasioned categorisation practices form a resource for the local organisation of knowledge in the situated viewing of a scene" (Smith 2021, 184). Just as the identification of "where we are" is "a categorial accomplishment" (Smith 2017, 121) for cyclists and motorists sharing the road, or as the runner "is not simply 'seeing' a world through which they are running but are, moment-by-moment, assembling the intelligibility and sense of that world, as a categorially organised, occasioned corpus" (Smith 2019, 38), the player in Geoguessr is constantly categorically assembling what they are seeing, even though they are doing so in and as of the course of a very different activity (playing a computer game) and often without talking.[2] The live work of geographical localisation, which is in the focus of my analysis in this paper, can be depicted as follows:

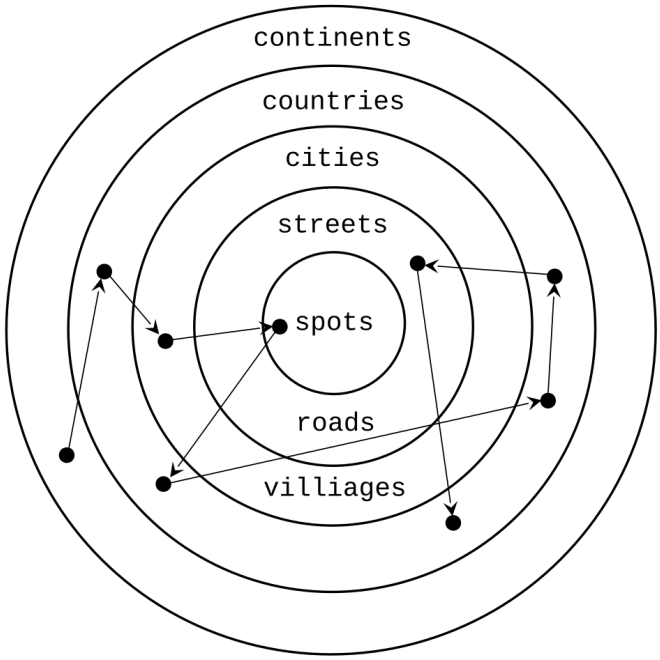

As the navigation progresses, the player jumps from one circle to another, without any predetermined order, responding to the new details brought to the fore by their situated activity. Sometimes there are ambiguities that can be temporarily tolerated, such as the equal possibility of a given spot being located in several different (often neighbouring) countries. Geo-

2 This limits—but does not undermine—the usefulness of studies of the "formulating place" in the talk (e.g., Schegloff 1972; Housley and Smith 2011; Kitzinger et al. 2013) to the present paper's concerns.

graphical categories are not "applied" but assembled according to the practical demands of the game. They form the *occasioned corpus*[3] used by the player to organise their activities. In the descriptive analysis to which I will turn next, I will try to trace the relationships between the unavoidably local, contingent work of "Gestalting", that is, foregrounding/backgrounding the setting's details elicited by player's searching actions, and the nested ordinary organisation of geographical knowledge.

ANALYSIS: NAVIGATION FOR, AND AS, CLUE-FINDING

In this section I will present the descriptive analysis of an episode from my own history of playing Geoguessr. It demonstrates the successful (maximally scored) finding of where I am in 3 minutes 45 seconds. The order of discussion will be chronological: I will follow the events of my search as they unfolded, making analytical observations about the particular phenomena that transpire in the process.

The search begins from the scene shown in figure 1. Let me reproduce it here to make it easier for the reader to find correspondences between the images inserted in the text and what I have to say.

This scene can tell different things to different observers. Here is what it tells to me, a relatively novice player.

The place looks like a city or a town. The size is hard to guess, but it could be large, judging by the motorway visible on the left. You can also see a sign and a text above what appears to be a garage door on the right. The characters are Latin, so this is probably a country where the main language is Latin-based. Yes, Latin characters can be found all over the world, but the type of the architecture, the type of the vegetation, and the type of the car I can see in the image point to a European country or the UK.

This is just a first glimpse of the setting, but it already shows the phenomenon that is massively relevant to Geoguessr: ordinary geographical knowledge is *expandable*. In the previous section I depicted a non-linear trajectory of moving between geographical categories when finding out where you are. This trajectory not only reflects the situated work of guessing, reevaluating, searching for clues, doubting, gaining insights, and so on, but also serves as learning curve for the player who follows it. As the geographical knowledge is central to playing Geoguessr, the results of the player's attempts to find where they are depend fundamentally on player's developing skills in assembling geographical categories for the practical purpose of recognising the scene they see on the screen.

3 The notion of "occasioned corpus" can be traced back to Harold Garfinkel's study of juror's work (1967, 104–15). Using Felix Kaufmann's (1958) concept of "corpus of knowledge", Garfinkel shows that jurors assemble their own "corpus of fact", "treated by the jurors at any given time as 'the case.' By 'the case' is meant the logical mode of 'actual' and is contrasted by jurors with the logical modes of 'supposed,' 'possible,' 'fanciful,' 'hypothetical,' and the like" (1967, 107). The preoccupation of the computer game players is certainly different from the preoccupation of the jurors, but the production and usage of the occasioned corpus is the core activity for both.

Figure 1

There are three ways of expanding geographical knowledge. First, the more you play the game, the easier it becomes to recognise certain places and to distinguish between them by the type of vegetation, the characteristics of the buildings, the peculiarities of the local writing systems, and many other things. In the course of the game the player not only memorises particular scenes, but also learns how to make this or that detail speak to them.

Second, the player can expand their geographical knowledge by acquiring knowledge outside the game. For example, they may surf Google Maps or read the geographical atlases and encyclopaedias in their spare time. More experienced players can make a guess based, for example, on the type of the road marking. In the image presented above, you can see yellow-painted street curbs and flagstones. This may be the type of marking specific to a particular country or the part of the world. You can expand such knowledge and thus make new details of the image come to the fore to tell you of what location they are details.

Finally, you can expand your geographical knowledge by expanding the "production knowledge" of Google Maps images. Consider these three snapshots from the video of the episode I am analysing here:

Figure 2a

Figure 2b

Figure 2c

In these images, I have marked with red circles the possible telling details of the images related to their production. Google Street View images are created by cars (and occasionally by people on bicycles or on foot) equipped with 3D cameras that constantly capture everything around them as they drive through selected locations. These cars have significant differences: they are of different colour, their equipment is different, and they are of different brands and models. Despite Google's efforts to erase their cars from Google Street View images, sometimes you can see some properties of them, such as type of antenna (short black stick in fig. 2a), colour (blurred grey-white spot in fig. 2b), or type of equipment (shadow in fig. 2c). All these details can tell you something about the location, because Google used different combinations of cars and equipment to 3D-capture specific parts of the given country. So, some users (not me) are able to guess where they are based on the type of antenna and the colour of the car.[4]

4 A lot of ingenious examples may be found in the "The Complete Guide to Russian Car and Seasonal Metas" created by Geoguessr user BarrBarrBinks: https://docs.google.com/document/d/1KMRVfyEDF5h-wYRkCCDVjaNh7XCy3xcbpTnRfybUoDMU/edit.

When I am talking about the expandability of geographical knowledge, I am not referring to the sheer amount of geographical information that the player knows, but to the readability of scenes in the course of the game. Even the first still image of the scene can tell a lot to the trained eye. Good players can identify the location with an error of several hundred kilometres being exposed to the image for 0,1 seconds.[5] But such efficacy does not reflect the scope of the information stored in the player's memory. Here, memory, like vision, is a practical achievement that can only operate on the basis of skilful work of inspecting the image for telling details and making those detail tell you something.

With these observations in hand, let us now return to the episode that we started to go through.

The initial image from which the player starts is inspectable in four ways: (1) you can simply examine what you see there, as we did above; (2) you can look around to see what else you are in the middle of; (3) you can zoom in/out on certain parts of the image; and (4) you can "move on" from the initial spot using the tools provided by Google Street View. In my case, after scrutinizing image for about 7 seconds, I zoom in on the car.

I've noticed a sticker on the boot lid of the car and decided to take a closer look. Vehicles of all kinds can provide valuable clues, as some of them are branded and contain information about the companies that own them, such as the full postal address, email address, website address, city, and so on. The vehicle registration plate can also be very helpful. Although Google blurs out all the number plates on cars, you can still sometimes make out their characteristics, such as the colour. Zooming in on the car, however, proves useless for me. The sticker contains no information (red cross on white background is too ambiguous to be helpful) and I don't know this type of the plate.

Figure 3

5 See how one of the most famous players in Geogeossr, Trevor Rainbolt, identifies 20 countries by the images being exposed to him for 0,1 seconds: https://www.youtube.com/watch?v=ff6E4mrUkBY.

This failure to turn a detail into a clue is quite common for the novice player in Geoguessr. Sometimes it takes a long time to find a clue: you might be moving down the road in some outdoor area for several minutes before you come across a road sign or a human settlement. And now and then the clues are too ambiguous or uncertain to be of immediate help. For example, often you may be able to find the name of the road or of the town you are in quite easily, but if this is not a motorway or a big city, you won't be able to use the name to find out where you are. Occasionally you can guess the country from the text on the road sign, but this is still too broad a category to be useful for finding your exact location on the map. Players learn to establish stable connections between the types of details that they can bring to the fore and the types of the clues into which they can turn those details. In my case, I zoomed in on the car to inspect the sticker, ignoring the colour of the car and its brand, because I knew that text on a sticker could be informative for my search. The car here seemed to be the most promising source of the clues of all the objects visible from the initial position.

So, what do I do now that the inspection of the car has brought nothing? Instead of looking around, I keep going. I see the road at some distance ahead of me and I jump towards it.

There are two types of "movement" in Google Street View, both shown in figure 4. You can take a "step" by clicking on the white arrow or pressing the arrow key on your keyboard; or you can point to a specific spot—in this case, indicated by the hand-shaped cursor inside the black circle—and "jump" there. The second method is quicker, so I use it more often. However, if you only need to move a few metres (for example, to get a better view of something), the first method is more efficient.

Why didn't I look around? In retrospect, it seems like I found the road in front of me to be the next most promising place. Roads, especially if they have some visible signage, like this one, are the main source of useful geographical information: there may be road signs with street names or road/motorway numbers, signs with the names of towns and villages, business signs on the nearby buildings, and many other relevant things. So, here I am not jumping just to get ahead—as sometimes happens in the game, for example, when you start in the middle of the forest or the desert—but I am jumping *for* something. At the same time, this "something" is not a particular object that, as I know, has a characteristic that has caught my eye. I don't know the exact details that this action will bring within my reach. I just expect that this new spot may contain *something* that can be turned into clues. Roads are hotspots imbued with makable-into-clues details.

Figure 4

Figure 5

The jump brings me to the new spot (fig. 5). Now I have new objects and spots available to me. From here, I can plan a further route amongst them, or I can look around for new clues about where I am. Here, one object immediately catches my eye, even without looking around: the road sign on the left. Road signs can contain a lot of clues, so I decide to have a closer look. I do this, however, not because I *know* that roads signs may be useful, but because I've made a *situated noticing* of it as a result of my detail-revealing jump. When I jumped to the new spot, the road sign jumped into my practical field[6] as an inspectable-for-clues object. I zoom in on the sign and here is what I see.

6 I draw the notion of "practical field" from Maurice Merleau-Ponty. In *Phenomenology of Perception* (2012, 84) he discusses the practical field as a totality of "manipulable objects" that appeal to the body of the actor.

Figure 6

I am in the Netherlands! I don't speak Dutch, but I've been to the Netherlands several times (the last time was recently) and I know some German which can be quite similar to Dutch. For example, here I can see "VERBODEN", which reminds me of the German word "verboten" ("forbidden"). Given the format of the sign on which the text is written, I can assume that it forbids something. The word "VOOR" in the text also looks like characteristically Dutch, with the double "o" being ubiquitous in public spaces in the Netherlands.

Now I know where I am. My "knowledge", however, refers to the category "countries", which is not enough for finding my exact location of the map. But it is still very helpful: being able to pinpoint the country of my initial position cuts off the outer circles in the nested scheme of ordinary geography with which I operate. There is no need any more to worry about "continents" or "parts of the world". I can focus now on the inner circles, the most important of which is a "city/town", because without it all inner-circle information becomes useless. There is a good chance that the next clue that I find will be the name of the street, but without the name of the city/town in which it is a street I cannot find myself on the map.

This means that I have to keep searching. I jump forward again along the short lane that I am "standing" on (the destination is shown in fig. 7), and when this doesn't bring anything new, I jump to the next spot (fig. 8 shows the target)—all to get closer to the road in front of me.

Figures 7 & 8

The second jump takes me to the area with several road signs from which I might be able to pick up some useful information. And indeed, when I get there, I zoom in on the post with two signs which usually contain street names.

Figure 9

I can read the street names clearly, but they tell me nothing, except to confirm that I am in the Netherlands. They may come in handy later, as street names are instrumental in finding the exact location, but at this stage of my search they can only be memorised for the future (or I can memorise where to find them in case I forget the names). These two street names are at the moment detached from some of the outer circles of my geographical knowledge: I cannot relate them to the "cities/towns" as don't know the name of the current city/town and I cannot infer its name from the street names. However, this detachment only reflects the current state of knowledge assembling that I am doing for the practical purpose of finding out where I am. Later findings should allow me to build the links between all the circles, without which I won't be able to pinpoint the precise location of the initial point in the game.

Another interesting feature of the street names is that they are of a low *location-specificity*. We can ask about every category of common-sense geography "where is it?", and we usually expect that the larger the category, the more location-specific it is: almost everyone knows "where" is Europe (it is on the planet Earth and is not in Africa), many people know "where" is Spain (it is in Europe), fewer people know "where" is Madrid (it is in Spain), and even fewer people know "where" is Manuel Luna Street (it is in Madrid). Geographical categories are location-specific in what they tell us about "where they are". As a rule, street names tell us nothing about where they are. Except for unique streets that are associated with a particular city/town, street names are often reproduced in difference cities/towns in the same country and even in the different countries that speak the same language.

What Geoguessr, as a geography computer game, adds to this distribution of location-specificity, is that to be successful in the game, the player has to put geographical categories

on the map. It is not enough to know that Madrid is in Spain. You have to know *where* in Spain it is. Therefore, there is a double structure of geographical categories in Geoguessr: the player needs to find out both their *logical* relationships and their *cartographic* relationships. And both are established in the course of playing the game, from within the unfolding sequence of the player's situated actions.

I have to move on. On the opposite side of the road there is another set of street signs (fig. 10), which becomes visible as you zoom out from the signs with street names. The visual organisation of the image produced by my actions foregrounds these signs and highlights their inspectability. However, I cannot read the signs from this spot, so I make two jumps to get closer to them, and zoom in on them (fig. 11).

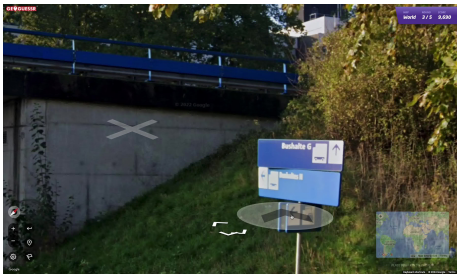

Figures 10 & 11

They tell nothing new. The language looks like Dutch again, but there is still no indication of the city. This is quite common with smaller roads: you can find information about their names (although even that can be problematic in some parts of the world), but you have to search for the name of the city/town and that can take quite a while. Major roads are more helpful in this respect, and as I happen to have one nearby, I decide to try my luck with it. The decision is contingent in the sense that, although I had noticed that there was something like a motorway near the place I am inspecting, it wasn't until I zoomed out from the set of road signs shown in figure 11 that I noticed that there were also large road signs above this motorway (fig. 12).

Figure 12

You could definitely see part of the back side of these motorway signs from the initial position (see fig. 1, on the left), but until this moment they didn't really exist for me. It is my actions that made them appear. In this respect, my actions are temporary vectors that structure the scene so that certain details of the images pop up along them. They are vectors because they project the next actions, and they are temporary because they are constantly changing in response to the details which transpire as a result of player's doings. But they are also vectors in the categorical space of ordinary geography, which is constantly being rearranged as the situated configuration of geographical knowledge assembled by the player.

So, I zoom in on the motorway signs from where I am located.

Figure 13

This is the maximum zoom, and I spend the next 18 seconds examining the signs. What do they tell me? Usually, these large signs do not tell you what city you are in, because they give directions to other places. But they can contain the names of major cities and the numbers of the nearby roads. Here we can see the red sign-in-sign "A12", the white sing-in-sign "RING", and the names of four cities: "Rotterdam", "Utrecht", "Den Haag" (I wasn't sure about "Den" due to the image resolution, but I could definitely see "Haag"), and "Delft".

This should mean that I am not in one of those cities. All the other cities/towns that belong to this category in the Netherlands are equal to me at the moment. They are not equal in terms of my knowledge of these cities, some of which I know better than others, but in terms of the collection of clues about the place that I've managed to pick up so far. All of them are equally possible candidates for being the city/town that surrounds me. As I accumulate information about the place, some differentiations between members of one category become more relevant. For example, based on the type of architecture I could see in the images, I might guess—as I did—that I am in a large city rather than a small one. In this case, I would start to group the cities/towns in the Netherlands into these two classes. But if I can only put the place where I am in into a *group* of places, I have not achieved my goal. I need to identify the *unique* city/town, i.e., I need to find this place as *this* place, and not as similar or equal to other places. The common features that *this* place share with other places are only important as the features of *this* place. This means that the geographical categories used in Geoguessr are put to work depending on the assembled sense of the place where the player finds and moves themselves.

The motorway signs narrow down the range of possibilities with regard to the location in which I am, but the range is still wide. In addition to the names of cities that can be excluded, another useful piece of information is the number of the motorway (A12), because it is easier to find numbered motorways on the map than streets. However, it seems that this is not the number of the motorway over which the signs are hung, but of some other motorway. I better look for something else. It's usually a good idea to get to the motorway to see what it has to offer, but I don't see any access to it nearby, so I zoom out, turn around, and start to move down the road.

I make three jumps before I notice the large sign on the wall of the multi-storey building ahead of me. I stop, adjust my perspective, and zoom in. This is what I see:

Figure 14

The black building looks like a hotel called "...den Tulip", probably "Golden Tulip". The Netherlands is famous for its tulips, so it makes sense. Unfortunately, there is no indication

of the city in which this "Golden Tulip" is located. Still, it's worth checking out the place, as there might be some additional information around the hotel, like more signage or cars with the names of local businesses (e.g., delivery service companies). So, I make another jump to the beginning of the lane that runs alongside the hotel, turn a little bit to the left, find nothing interesting there, then turn to the right and notice the sign on the right which I immediately zoom in on.

Figure 15

There are the name of the hotel and, below it, the name of the city. It's The Hague ("Den Haag")! This finding makes all the other pieces of information that I can collect inspecting this scene irrelevant. I do not even notice the name of the café/restaurant on the ground floor of the hotel or the blue sign with white arrow on the left. The name of the city dominates my visual field because it is the point of convergence of logical and cartographic relationships between geographical categories. On the one hand, I now know where I am in the Netherlands (I am in The Hague). On the other hand, although I don't know where in the Netherlands The Hague is and where in The Hague I am, the name of the country and the name of the city allow me to turn to the map and to try to locate the city. The relatively small size of the country and the relatively large size of the city (the fact that I know the name of the city without knowing much about the Netherlands speaks itself to categorical membership of the city) are going to be helpful in this, due to the scale organisation of Google Maps: I can find the continent, *then* the country, *then* the city.

Here, as in the previous parts of my search, I use the geographical information in and as a material organisation of the things visible in the images. Geographical knowledge is not only something that we can use in talking and writing, but it is also something visible and findable. When we move in the real world, we move amongst road signs, signboards, plates with numbers of the houses, signs over the entrances with the names of the shops, mileage signs, and so on. Players in Geoguessr use all this information when searching for clues. And they know where to find it. For example, if you want to find the name of the street, you better go to the nearest crossroad where you will usually find signs with street names. The applicabil-

ity of the categories of common-sense geography to both the real world and the pictures of the real world is based on this materiality. But the materiality also includes maps (printed or electronic) where the geographical information is presented in the form detached from the actual setting and has to be brought into correspondence with it.

The material organisation of geographical knowledge that players assemble can play tricks on them. Consider the signboard by the "Golden Tulip". The fact that it is the stationary signboard is not insignificant. Unlike moving objects, such as cars, which can carry signs and texts that have nothing to do with the place where they are encountered, stationary signboards are more reliable. To be sure, there are some clues that can be gleaned from the texts on cars (for example, taxi companies often operate in certain cities; or, if you see several differently branded trucks with the same city indicated on them, chances are that this is the city where you are in at the moment). On the other hand, stationary signboards can also be deceptive. For example, I once saw several pharmacies called "Farmacia Guadalajara" in the city and was sure that I was in the Mexican city of Guadalajara, only to be surprised to find out that this is the chain pharmacy that can be found throughout Mexico and the actual city is Tampico. Here the situation seems clearer: it's hard to read the bottom line on the signboard in any other way. But we must also remember that the certainty of identification can be undermined at any moment. After examining the road signs on the motorway, I was sure that this was not The Hague. Now I am confident of the opposite.

In any case, it looks as if I now have a decisive clue. It is still not the exact location of the initial position, but knowing the city significantly reduces the area you need to inspect to find where you are on the map. The Hague is still a large city. Although I don't know how large it is—and I don't know how to define the "largeness" of a city—I think it's smaller than, say, Paris or Berlin. Also, on the signboard there is a word "Zoetermeer" before "Den Haag". I don't know what it means. I've never been to The Hague and I don't know its geography. Is it the name of a street? Could be, although, if this is an address, I would expect the street name to be accompanied by the number of the building. Is it the name of a district? Also possible, and this would be an even more reasonable guess.

Figure 16

At this point, I turn to the map for the first time. I click on the map window in the bottom right corner of the screen to enlarge it so that I can navigate around it. Why is the map brought up only now? The answer is that there was nothing to look at on the map. Sure, the map is full of stuff, but it would not help me. I could open the map of the Netherlands right after finding out that this is the Netherlands, but it would tell me nothing. Now the situation is different. If I find The Hague, I can inspect it. I know that scale matters in Google Maps: the smaller the scale, the more geographical detail you see. But you also reduce the size of the inspectable area. So, I open the map not to find the exact location (I don't even know which street I am on), but to familiarise myself with the city and its geography. It can tell me something (the name of the street, the number of the motorway, etc.) related to what I already know about the place. For this purpose, it is enough for me to have a bird's-eye view of the city.

Finding The Hague on the map is not as trivial as it might seem, because I have no idea where it actually is in the Netherlands. What also hinders navigating the map is that some city names only appear at a certain scale (see fig. 17, where I zoomed in and got several Dutch cities, but not The Hague). To get this level of granularity, you have to zoom in, but you will see the smaller part of the map. Luckily for me, the Netherlands is a relatively small country, so I zoom in one more time and get the city I am looking for (fig. 18).

Figures 17 & 18

What we can see here is the consequence of the specific way in which Google Maps visualises geographical categories: a given scale only brings up the cities of a certain size. The basic rule is that members of the same category are displayed in the same way. The commonsense geographical categories may or may not match those of Google Maps (for example, I may group as "large" the cities that the map would consider to be of different sizes), but using the map means adjusting your search depending on what the map offers you. The category is an instrument: at the end of the day, I am looking for a particular city, and not for a categorically defined set of the cities such as "big cities", or "capitals", or "small towns". But I know that to find the particular city, I need to find it among all the cities of the same type that the map shows me.

Having found my target, I zoom in further, and the next time I zoom in I get this:

Figure 19

The first thing I notice on the map is "Zoetermeer" to the right of "Den Haag". The same "Zoetermeer" that I can see in the picture in front of me. At this point I am sure that Zoetermeer is a district of The Hague, not the least because it is in the same line as "Den Haag" on the hotel signboard and because of its proximity to The Hague on the map. I completely inattentive to the fact that its name on the map is in bold, which, according to Google Maps navigational notation, means that it's a separate town. I figured this out much later, when I googled "Zoetermeer" a few days after the game. Here I miscategorised the place, but it turned out to be of no real consequence to my search, because I have found the place on the map *before* I have started looking for it among the other members of the same category. In other cases, however, it is the other way round, and sometimes I am laboriously—and often in vain—searching for a particular place among all the places of the same size that I get on the map at a given scale.

Now I zoom in on Zoetermeer to see what comes up. I still have no hard clues, i.e., something particular to look for. It is possible, as has happened many times before, that after zooming in on Zoetermeer and inspecting the map, I will get nothing and will have to go back to Street View and search for other clues. I zoom in several times until I get to a scale small enough to see some geographical names but not too small to get swamped with the details of the tiny spot (fig. 20), and start moving the map to one side and then to the other, zooming in a bit more. Suddenly something very interesting appears at the bottom edge of the map (fig. 21).

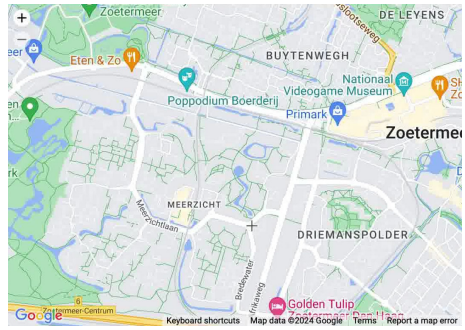

Figures 20 & 21

It's the Golden Tulip! It is just a happenstance, of course, I was not looking for it, but my previous encounter with this name and the sequence of my actions after that encounter make the name immediately recognizable on the map. What I also know from my previous search is that the initial position where I was dropped is not far from the hotel and I remember the pathways between the two. The Golden Tulip can serve as a reference point when finding the precise location of the start. But to do this, I need to establish more definitive connections between the map and the Google Street View images. At the moment they are linked by the name of the hotel. And its categorical membership is not unimportant for my search. I have no doubt that this is the same Golden Tulip that I can see in Google Street View, although I know that there can be several hotels with the same name even in the same city. But I also know that hotels, as a category, are not as numerous as, say, chain stores, so the chances of having several hotels with the same name in such a small place as Zoetermeer are very low. Moreover, I have a way of checking my guess: I can zoom in on the Golden Tulip to see if what the map shows corresponds to what I can see in the picture. But I don't need to, because I am sure of my identification and because I can do this as a part of my search for the solution.

So, now I can start to pinpoint the precise location of the initial spot. I zoom in on the hotel on the map (I stop when I can see the roads around it) and click on the "Return to Start" button.

Figure 22

Zooming in on the map gives me the details that allow me to correlate the map with the Google Street View images. Besides the Golden Tulip, the most helpful part of the map is the layout of the roads around the hotel. The roads on the map look very different from what you see in the image, and sometimes finding out exact location can be tedious. But I remember that to get to the hotel I need to go forward and then turn right, and this can serve as a general clue for finding correspondences between the map and the images and for finding the starting point. In my case it went like this: I zoomed in on the map (fig. 23a), jumped forward to check out the U-shaped road that is very vivid on the map (fig. 23b), turned right to see if the building on the right matched the dumbbell-shaped building seen below the U-shaped road on the map (fig. 23c), then moved forward a bit to make sure that this was indeed the U-shaped road (fig. 23d), then returned to the initial position (fig. 23e), turned right to calibrate my position relative to the building (fig. 23f), turned back to calibrate my position relative to the end of the cul-de-sac behind me (fig. 23g), put the pin on the map, and then clicked "GUESS".

Figures 23a & 23b




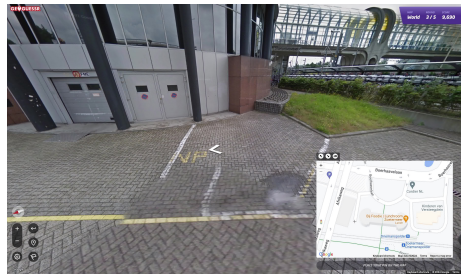
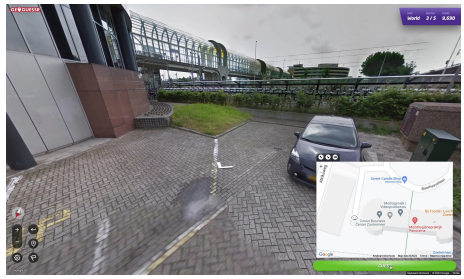

Figures 23c–23g

All for the result:

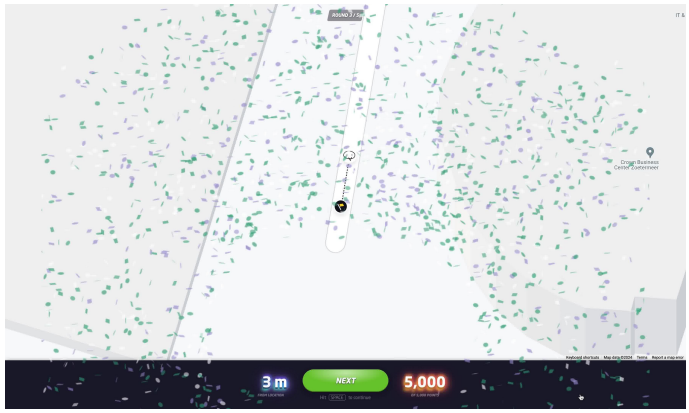

Figure 24

It's worth noting that the point where I put the pin is located both in the logical space of relationships between geographical categories and in the cartographic space of roads, streets, buildings, etc. It is on the particular street that has a name and physical properties. However, when I was finding the precise location of the starting point at this final step, I was finding only its cartographic location. Its logical place among the geographical categories has already been determined. After finishing this round of the game, I could not tell on which street the initial spot was or even in which part of The Hague it was. (As I said, actually it wasn't even in The Hague.) To find the solution at this final step, I had to establish the correspondence between the map and the Google Street View. The already assembled configuration of categorical knowledge was enough for the practical purpose of initiating and carrying out the cartographic work. And this cartographic work was done by manipulating the images and the map to foreground the details that could be brought into correspondence with each other.

The presented descriptive analysis shows that, practically speaking, playing Geoguessr consists in assembling the occasioned corpus of geographical knowledge. This is done by contingent Gestalting, foregrounding/backgrounding, of the details of the images and of the map through the sequences of situated actions that respond to and elicit the emerging features of the setting. The player moves simultaneously in two planes: logical and cartographic, connecting what they see to the logical relationships between geographical categories and to the unfolding succession of scenic properties. Searching for clues being the main preoccupation of the player, they constantly answer the question "Where is it?" (or "Where am I?", which is pragmatically identical in the case of Geoguessr) by detailing the images of the real world and the map using the available hardware and software resources, as well as their own bodily resources. "Where it is" can be determined with different levels of granularity, but if the player wants to score the maximum for the round, they need to assemble the occasioned corpus that can lead them to the precise location on the map. In the episode presented, I could put the pin "somewhere in the Netherlands" early on and I would get a decent number of points. But to the get the maximum score, you have to try to get to the heart of the scheme I draw in the second section of the paper, i.e., to the "spot", and to do that you have to work your way up by constantly guessing the categorical membership of what you see and narrowing it down to the particular place.

Now, to address to the topic of this special issue, I would like to discuss the prospect of describing playing Geoguessr as a discovery practice.

DISCUSSION: IS PLAYING GEOGUESSR A DISCOVERY PRACTICE?

It is always possible to reduce the discussion of discovery practices to the matter of ordinary language. In this respect, the answer to the question posed in the title of this section is certainly "no". The Age of Discovery is over, and no one today would call the search for a geographical location a "discovery". There are still situations when it is justified to talk about

discovering a place. For example, you might say that you have discovered a nice vintage shop. But the process of finding the shop would not be described as "making a discovery".

Another aspect of the ordinary usage of the word "discovery" is that it is strongly associated with science. "Scientific discovery" has become a shibboleth, almost synonymous with "discovery", so that to call a practice a "discovery" practice inevitably invites the comparison of that practice with the science. In this respect, too, playing Geoguessr is not discovering. There is nothing scientific about playing a computer game. (And scientists are no better computer game players than anyone else.)

But if we focus on the structures of practice instead of ordinary language, the answer to the question is no longer that straightforward. Finding one's place in Geoguessr and discovering the "hidden gem" of the nice little shop have, at their core, procedures, albeit different ones, of making something previously unknown known. In the second case, the procedure implies the feeling of surprise that is absent in the first one, and in the first case the actor is aware of the unknown's existence and character, unlike in the second case, but in both cases the actor acquires geographical knowledge about the particular place—"where it is".

On the other hand, playing Geoguessr has some features in common with scientific discovery. Although what the player in Geoguessr finds out is, to use Kenneth A. Strike's terminology (1975, 466–467), more akin to "relative discovery" than "absolute discovery"[7] that is meant by "discovery" in science, the player in Geoguessr follows the similar path from "an evidently-vague IT" to the "'relatively finished object'" (Garfinkel, Lynch, and Livingston 1981, 135) that scientists follow. The objects of their search are obviously different, with the players looking for the cartographic location and the scientists looking for "The Independent Galilean Object" (to paraphrase Garfinkel, Lynch, and Livingston's term "The Independent Galilean Pulsar"; see 1981, 138). But what they ultimately find is in both cases the feature of their situated work of finding it.

These remarks are not intended to equate playing Geoguessr with other discovery practices, whether mundane or scientific. Some similarities sit alongside multiple differences. However, I believe that we can consider playing Geoguessr *together* with discovery practices, even if we cannot call it "discovery practice" itself. I would suggest applying to them Wittgenstein's famous concept of "family resemblance" (Wittgenstein 2009, 36e), with one amendment: as in any family, there are close relatives and distant relatives. Among the discovery practices, playing Geoguessr has only a distant resemblance to "first time through" scientific discoveries or accidental mundane discoveries. Other discovery practices are closer.

---

7 "Absolute discovery requires

1. that *X* found out *p* for himself [*sic*], and
2. that *X* is the only discoverer of *p*.

It will also require either

3. that *X* was the first to know that *p*, or
4. that *X* was the first to successfully transmit that *p* to the world" (Strike 1975, 466). The "relative discovery" is the discovery where only criterion (1) is applicable. This means that learners can make discoveries for themselves without discovering something for everybody.

The closest is probably the work of police investigation, which has a lot in common with playing Geoguessr. The player combines two roles: the role of the crime scene investigator (Williams 2007) and role of the investigating police officer (Sanders 1977). Players search the environment for clues that might help them determine where they are, just as detectives (both real and fictional) collect and examine the available evidence to determine what happened. This journey from what the scene presents to the final clarity demonstrates similar phenomena in both cases: guessing, finding yourself wrong, connecting pieces of information, overlooking some details, and so on. Both consists in "solving the mystery." Of course, there are also fundamental differences. Detective work, for example, implies the relevance of the "incongruity procedure" (Sacks 1972, 283), i.e., of the suspicion that the appearance of the person or object hides something, that it is not what it seems to be. In Geoguessr, although the player sometimes doubts their identification of the scenic details, there is nothing like suspicious that things are not what they seem to be. Rather, the player presumes that they haven't seen enough or that what they have seen is too vague to serve as a clue.

That said, I would argue that playing Geoguessr belongs to the family of "discovery practices". I think it has too much in common with some of the discovery practices. And although it is impossible to list the properties of discovery practices "in general" (which is precisely why Wittgenstein introduced the concept of "family resemblance" between practices, in his case games), the examinable similarities between playing Geoguessr and, say, detective work can be helpful in understanding why both might be attractive in their own way to computer game players and people considering a career in the police or reading a detective story.

CONCLUSION: WHAT HAVE WE LEARNED?

In this paper I have considered several key features of playing the geography computer game Geoguessr as a practical activity. My focus was not on the properties of the game *as* a computer game, but rather on the ways in which players are engaged in the work of finding out their location in 3D representation of the real world. Although this work has much in common with other computer games, I was primarily interested in understanding the characteristics of playing the game that it shares with other discovery practices, more habitual to everyday-life and scientific inquiries. For this reason, I didn't situate my analysis among other studies of computer games, but instead highlighted the differences and similarities between the practical organisation of playing Geoguessr and the organisational phenomena found in other discovery practices. The practical organisation that I have attempted to describe involves assembling the occasioned corpus of geographical knowledge over the course of the game by manipulating the details of the scenes on the screen in search of the clues to answer the main question: "Where am I?" (or, alternatively, "Where is it?"). I have tried to be as faithful as possible to the practice of playing the game, but there is obviously still much to be improved in this respect. But whatever the shortcomings of the analysis, I hope that I have been able to provide sufficient grounds for serious consideration of where else we can

find phenomena similar to those of "classical" discovery practices, and perhaps to motivate the readers to play Geoguessr themselves in order to continue, correct, or reject the analysis presented here.

REFERENCES

Benford, Steve, Andy Crabtree, Martin Flintham, et al. 2006. "Can You See Me Now?" *ACM Transactions on Computer-Human Interaction* 13 (1): 100–33.

Garfinkel, Harold, Michael Lynch, and Eric Livingston. 1981. "The Work of a Discovering Science Construed with Materials from the Optically Discovered Pulsar." *Philosophy of the Social Sciences* 11 (2): 131–58.

Gurwitsch, Aron. 2010. *The Collected Works of Aron Gurwitsch* (1901–1973), Volume III: *The Field of Consciousness: Phenomenology of Theme, Thematic Field, and Marginal Consciousness*. Springer.

Housley, William, and Robin James Smith. 2011. "Mundane Reason, Membership Categorization Practices and the Everyday Ontology of Space and Place in Interview Talk." *Qualitative Research* 11 (6): 698–715.

Kitzinger, Celia, Gene H. Lerner, Jörg Zinken, Sue Wilkinson, Heidi Kevoe-Feldman, and Sonja Ellis. 2013. "Reformulating Place." *Journal of Pragmatics* 55: 43–50.

Laurier, Eric. 2001. "Why People Say Where They Are During Mobile Phone Calls." *Environment and Planning D: Society and Space* 19 (4): 485–504.

Laurier, Eric, and Barry Brown. 2008. "Rotating Maps and Readers: Praxiological Aspects of Alignment and Orientation." *Transactions of the Institute of British Geographers* 33 (2): 201–16.

Laurier, Eric, Barry Brown, and Moira McGregor. 2016. "Mediated Pedestrian Mobility: Walking and the Map App." *Mobilities* 11 (1): 117–34.

Lewis, David H. 1994. *We, the Navigators: The Ancient Art of Landfinding in the Pacific*. University of Hawaii Press.

Lewis, David H., and Mimi George. 1991. "Hunters and Herders: Chukchi and Siberian Eskimo Navigation Across Snow and Frozen Sea." *Journal of Navigation* 44 (1): 1–10.

Merleau-Ponty, Maurice. 2012. *Phenomenology of Perception*. Routledge.

Pehkonen, Samu, Thomas Aneurin Smith, and Robin James Smith. 2022. "Maps, Mobility, and Perspective: Remarks on Map Use in Producing an Orienteering Course." *Mobilities* 17 (1): 152–78.

Sacks, Harvey. 1972. "Notes on Police Assessment of Moral Character." In *Studies in Social Interaction*, edited by David Sudnow. The Free Press.

Sanders, William B. 1977. *Detective Work: A Study of Criminal Investigations*. The Free Press.

Schegloff, Emanuel A. 1972. "Notes on a Conversational Practice: Formulating Place." In *Studies in Social Interaction*, edited by David Sudnow. The Free Press.

Smith, Robin James. 2017. "Membership Categorisation, Category-Relevant Spaces, and Perception-in-Action: The Case of Disputes Between Cyclists and Drivers." *Journal of Pragmatics* 118: 120–33.

Smith, Robin James. 2019. "Visually Available Order, Categorisation Practices, and Perception-in-Action: A Running Commentary." *Visual Studies* 34 (1): 28–40.

Smith, Robin James. 2021. "Categorisation Practices, Place, and Perception: Doing Incongruities and the Commonplace Scene as 'Assembled Activity'". In *On Sacks: Methodology, Materials, and Inspirations*, edited by Robin James Smith, Richard Fitzgerald, and William Housley. Routledge.

Smith, Thomas Aneurin, Eric Laurier, Stuart Reeves, and Ria Ann Dunkley. 2020. "'Off the Beaten Map': Navigating with Digital Maps on Moorland." *Transactions of the Institute of British Geographers* 45 (1): 223–40.

Strike, Kenneth A. 1975. "The Logic of Learning by Discovery." *Review of Educational Research* 45 (3): 461–83.

Williams, Robin. 2007. "The 'Problem of Dust': Forensic Investigation as Practical Action." In *Orders of Ordinary Action: Respecifying Sociological Knowledge*, edited by David Francis and Stephen Hester. Ashgate.

Wittgenstein, Ludwig. 2009. *Philosophical Investigations*. Wiley-Blackwell.