

Making Space: The Natural, Cultural, Cognitive and Social Niches of Human Activity

Barry Smith

Uncorrected proof version of paper now published in *Cognitive Processing* 22 (supplementary issue 1):77-87 (2021)

This paper is in two parts. Part 1 examines the phenomenon of *making space* as a process involving one or other kind of legal decision-making, for example when a state authority authorizes the creation of a new highway along a certain route or the creation of a new park in a certain location. In cases such as this a new abstract spatial entity comes into existence – the route, the area set aside for the park – followed only later by concordant changes in physical reality. In Part 2 we show that features identified in studying this phenomenon of legal spacemaking can be detected in other spheres of human activity, for example in planning (where spacemaking is projected into the future), and in reasoning about history (where spacemaking is projected back through time). We shall see that these features display themselves in especially complex ways in our everyday use of language, and we conclude by examining the implications of this complexity for attempts to create an artificial intelligence that would enjoy a mastery of language that would be equivalent to that of human beings.

Keywords: fiat objects, spatial ontology, boundaries, artificial intelligence

1. On fiat spatial entities

All organisms require space to satisfy practically all of their needs. In the case of our early human ancestors these needs were met by finding places which could be modified by physical means to make them suitable for habitation. In the course of time, however, space began to be parceled out in more abstract ways, so that in around 1400 BCE in Egypt we encounter what may be the first use of surveying practices to divide land into plots (in this case for purposes of taxation). Somewhat later,

we see plans and maps being used to mark out spaces in advance of the construction of buildings, for example the Plan of St. Gall from around 810–820, which is also one of the earliest preserved examples of an architectural blueprint (Froehlich 1981).

On April 30, 1883 an Act was passed by the State of New York establishing a state reservation comprising ‘certain lands in the village of Niagara’, with the goal of ‘preserving the scenery of the falls of Niagara’.¹ This Act resulted in various sorts of changes in the world, including physical changes such as the posting of signs and the creation of shelters and pathways. But the immediate effect of the Act was to bring into being new legal entities, including not only the site of the park, but also a new organization called the Niagara Parks Commission and a new role of Park Superintendent. Such legal entities are needed in advance of the building of fences, because a determination is required as to where the fences shall be positioned and who should be responsible for their construction.

The park and its boundary are examples of what we shall call *fiat spatial entities*, which are entities that are brought into existence through fiat demarcations of space, as achieved, for example by legal authorities drawing lines on maps.² Fiat spatial entities span the boundary between geometry on the one hand and the material reality of (in most of the cases relevant here) the Earth’s surface on the other.³

There are also *fiat spatiotemporal* entities, creating on the one hand what we can think of as tunnels through history as it has unfolded at the Earth’s surface, and on the other hand tunnels through future time, as for example the projected trajectory of a pilgrimage to Santiago de Compostela along the Way of Saint James, or of the 2024 NASA interplanetary mission to Europa, one of the Galilean moons of Jupiter.

2. Definitions and examples

We define a legal fiat spatial entity as the immediate result of a legally authorized demarcation of a certain region of space. Here ‘legal’ should be interpreted broadly enough that it will apply also to the results of such demarcation by any individual or organization which has the authority to make and realize plans involving demarcation of regions of space, as when a company picks out the site for a new factory, or when a church decides on the launch area for a new charity drive.

Fiat spatial entities may arise in many different ways. An aircraft performing an aerial survey of

¹ https://www.google.com/books/edition/General_Statutes/IYJZAAAAYAAJ?hl=en&gbpv=1&pg=PA155

² See Smith (1995, 2001).

³ The geometric aspect is illustrated most conspicuously by the spacemaking work of colonial governments in Africa, Australia and the Americas, where many of the created boundaries are in whole or in part rectilinear.

a given area might take thousands of photographs, each demarcating some rectangular fiat spatial entity. The whole series then serves as the first step in identifying portions of this area which might be set aside for specific agricultural purposes. Alternative plans for such a demarcation of the area might then be submitted for review and approval by a legal authority. These plans then contain maps depicting demarcations which bring into existence a new set of fiat spatial entities, which are in some sense experimental, and therefore not (or not yet) *legal* fiat spatial entities.

The family of full-fledged legal spatial entities includes not only the Niagara Falls State Park but also Canada and the United States, as well as the Province of Ontario, the State of New York, and all the other plotted⁴ countries, provinces, states, counties, cities, towns; school, postal and electoral districts; international time zones, urban planning zones, industrial and agricultural zones; the Schengen Zone, the Eurozone, the European Free Trade Area, the Black Sea Economic Cooperation Zone; as well as all the individual lots, tracts, and parcels of land which have been demarcated in legal enactments, title deeds, architectural blueprints and property developers plats over the centuries. It includes also all the archdioceses, provinces, dioceses, archdeaconries, parishes, and their counterparts in other geographically organized religions around the world. Someone who is standing inside any one of the above-mentioned circumscribed areas is entitled to use the word 'here' in such a way that this 'here' will be understood as referring unproblematically to the entire area.

Many of the mentioned entities are organized in hierarchies, with larger zones partitioned into smaller zones, which are in turn partitioned into still smaller zones. This reflects not only the hierarchical structure of the institutions of law, government, religion, and commerce, but also the reliance for almost all of human history on spatial proximity in effectuating management control. Organized religions, for example, as they established themselves in successive new territories, needed to ensure that the latter were organized after the pattern of their existing territories in order that they could be incorporated into the common hierarchical system.⁵

3. Fiat spatial entity creation through joining, splitting, and regimenting

Some of the changes in the legal spatial realm are a matter of the fiat joining or splitting of already existing legal spatial entities. Examples pertaining to entire countries are: Italy in the 19th century;

⁴ To describe a spatial entity as 'plotted' is to assert that it owes its existence to some demarcation typically effected with the aid of a map, as contrasted with, for example, the moon, which is set apart from its surroundings by *bona fide* physical discontinuities. Even small island nations such as Malta or Cuba are plotted in this sense, because they encompass smaller islands whose belongingness to the country in question is a result of demarcation.

⁵ In the Church of England, for example, there are two ecclesiastical provinces, headed by the Archbishops of Canterbury and York. England is divided into 43 dioceses, but there is in addition one further diocese, the Diocese of Gibraltar in Europe, one of whose component parts is the Eastern Archeaconry of the Church of England, covering Greece, Turkey and the former Soviet Zone.

Germany and Czechoslovakia in the 20th century; and all the countries of Europe, on a number of different cross-cutting organizational levels, in the 21st century.⁶ Some are a matter of the *redrawing* of boundaries, a process which involves the reassignment of portions of (or sometimes the whole of) one country to one or more other countries, as applied for example, to Poland, whose existence was suspended first between 1795 and 1918 and then again between 1939 and 1945.

The historically most important vehicle for the creation of new fiat spatial entities, however, consists in the making precise of boundaries left indeterminate by prior history or prior mapmaking, or which have otherwise become problematic. This may be because of newly acquired knowledge about or because of changes in the physical reference markers used to determine where the boundary lies, or it may be because of ambiguities or inconsistencies in relevant charts or legal documents. What we see here is the realization of an important function of the law, namely to promote regimentation in the realm of social interaction by making something formal out of what had previously existed in a more or less indeterminate and informal way. Laws enacted in order to fulfil this function turn what was previously a matter of loose expectations and rules of thumb into common, publicly defined and consistently applicable laws.

When changes in the realm of legal fiat spatial entities occur, they are of course associated with physical changes both prior to and in the aftermath of a legal enactment. An example of the former is the role of the Chernobyl nuclear disaster in bringing about the establishment of the Chernobyl Zone of Alienation, initially defined as a circular area of 30 km radius around the power plant to be placed under Soviet military control. An example of the latter are the physical changes involved in the dismantling of the Inner German Border when the German Democratic and Federal Republics were reunited in 1990.

Fiat spatial entities are abstract, in the sense that they have no material parts. As we saw, their creation goes hand in hand with the creation of other types of similarly abstract entities, such as roles, rights, laws, bylaws, regulations, contracts, tax obligations, and many more. The instances of each of these types are, to be sure, always in some way tied to matter. A fiat spatial entity is tied (for example) to the Earth's surface; a right is tied to the material entities (persons) who are the holders or grantors of the right. But the right is not itself made of matter, and it is not subject to the sorts of continuous causal changes to which all material entities are subject. In course of time, however, the *holders* of roles, rights, obligations and so forth will almost certainly be changed physically, as for example the newly appointed park rangers will acquire new capabilities as they experience the kinds of behaviour

⁶ <https://hofrat.ch/wp-content/uploads/2017/10/Supranational-European-Bodies.png>

appropriate to patrolling the park. Their legal authority will therefore gradually be complemented by a practical capability, as abstract rules become augmented by appropriate sorts of expertise (Dreyfus and Dreyfus 2005).

4. Reinach's ontology of the social act

The phenomenon of social entity creation was recognized already in 1964 by Searle in his work on how speech acts create obligations, and his account of this phenomenon was elaborated in greater detail in his later work on the construction of social reality (Searle 1995, 2010). But the first to explore the phenomenon in a rigorous manner was Adolf Reinach, a German legal philosopher who published in 1913 a monograph on what he called *The A Priori Foundations of the Civil Law* (Reinach 2013). Reinach uses the term 'social act' for what in the Anglo-Saxon tradition are called 'speech acts', both terms referring in effect to the same sets of instances, namely to those cases where human beings use language (understood as including moans, cries, gestures and so forth, including gestures involving pointing to lines on a map) in social interactions. Both Searle and Reinach took particular interest in what might be called normative social acts such as acts of promising, whose performance gives rise immediately to a claim and obligation on the part of promisee and promiser (Mulligan 1987).

Reinach defines a social act more precisely as an act performed by a human being to bring about some end by influencing another human being, typically by using sentences containing verbs such as 'promise' or 'order' or 'request'. A social act may involve direct interaction between one person and another, but many social acts achieve their ends indirectly through use of documents. Legal enactments⁷ are one example of social acts of this sort.

Reinach differs from Searle most importantly in that he develops a systematic ontology of legal and other social entities. Indeed Searle, for all his influential efforts to advance the ontology of social reality, went out of his way to dismiss the idea that there are *sui generis* social entities or social objects (Smith and Searle 2003, p. 302). This is because for him all that exists belongs to the realm of physics. A speech act is a special sort of patterned audio-acoustic blast; a cathedral is a special sort of patterned heap of stones. Talk of 'countries' and other legal spatial entities is in his eyes shorthand for talk about what he calls the 'deontic powers' which are imputed to the human beings involved (roughly, normative powers such as: to vote, to rule, to issue passports, to pass laws, and so on). What is of primary importance is then for him the way in which humans and groups of humans impute such powers to each other and to material entities such as border posts or passports.

⁷ Searle called them 'standing declarations' (2010, pp. 10-11).

Searle refers at one point to the idea 'that there is an element of magic, a conjuring trick, a sleight of hand' involved in the creation of, for instance, money out of pieces of paper and ink (Searle 1995, p. 45).⁸ In *Making the Social World*, indeed, he is ready to accept the conclusion that

money and other such instruments ... are products of massive fantasy. As long as everyone shares the fantasy and has confidence in it, the system will work just fine. But when some of the fantasies cease to be believable, as happened with the subprime mortgage instruments, then the whole system begins to unravel. (Searle 2010, p. 201)

Reinach, in contrast, accepts debts, contracts, claims, obligations, and so forth, as full-fledged denizens of reality. They are entities which are simultaneously both abstract and historical (a combination not recognized by philosophers hitherto). Such entities are tied together in structural networks with entities of other, complementary types. The *relation of ownership* between you and a certain *plot of land*, for example, is tied to the *title deed* that is currently in the *possession* of your *bank*, where it provides *secured collateral* for your *mortgage*.

All of these entities are brought into being by different sorts of actions on the parts of human beings, and they play roles in further such actions as time progresses, leading to the creation of further entities of the listed sorts. For example, the action of building a house on a plot of land will require, in most *jurisdictions*, a certain kind of *permission*, whose creation, again, will require a corresponding kind of legal (administrative) *enactment*, whose *ratification* will rest in turn on a certain kind of *legal authority*.

The mentioned created entities are abstract not only in that they are not made of matter but also in that many of them have properties of the sort we find in the realm of mathematics; a debt, for example, can be split into 100 equal parts, and sold off (as a right to repayment) to 100 different buyers.⁹ Fiat spatial entities, too, are abstract in the sense that they begin to exist as a result of non-physical processes and are thereafter subject to processes not physical in nature, for example of the sort through which a large parcel of land can be subdivided *on paper* into 100 separate plots.

The created entities we have been considering exist in every case on two levels: as general type and as particular instance. For example, the type *parcel of real estate*, as instantiated today in England, evolved from a state of affairs in the 11th century in which it had just one instance – the King owned all the land. From there it burgeoned forth in subsequent centuries to form the contemporary constellation of gigantic numbers of types and instances of property rights and parcels of real estate, reflecting the complex series of granting's of property rights and consequent subdivisions which have been effected as the law has evolved and been applied in different ways.

⁸ On Searle and verbal magic, see (Fónagy 2001, pp. 270 ff.)

⁹ On the economic consequences of such mathematical features of debt instruments see (de Soto 2000) and (Smith 2008).

5. The ambiguity of place names

To understand clearly what is involved here, it is important to note that the everyday expressions we use when referring to spatial entities in general and to places in particular are in multiple senses ambiguous. Consider the term 'Manhattan' as used in:

1. 'I live in Manhattan'
2. 'I live on Manhattan'

Sentence 1. can itself be understood in two senses:

In sense **1a.** it means: my home address lies in the Borough of Manhattan, the result of parceling out the area of New York City into five boroughs, which has the two-dimensional boundary that we see on an official map, and which came into existence on January 1, 1898 under the terms of the Greater New York Charter¹⁰. In this sense Manhattan has no physical properties other than its location. It is, in our terms, a fiat spatial entity.

In sense **1b.** it means: my home is one of the many sites which exist in Manhattan (as defined under **1a**), where 'site' is used to refer to what in biology is called the environment or habitat of an organism or group of organisms, also called a 'spatial niche'. A site is the kind of entity in which an organism or population of organisms can live (as bees can live in a hive, or bears can live in a cave); it is an open container in some surrounding material retainer (Smith and Varzi 1999, 2002; Bennett 2010). It is at the same time what offers affordances for the life activities of an animal (Gibson 1966). For dwellings in Manhattan it is the walls, floors and ceilings of the buildings which form the rooms in which we live and work, and it is the solid surfaces under our feet on streets and in parks which form the open spaces through which we move. Manhattan itself, under sense **1b** is a large assemblage of interconnected open containers; or in other words it is a very complex spatial niche.

Sentence **2** means: the building in which I live is built on and now forms part of Manhattan Island, an entity that is made of rock and concrete, soil, grass, trees, and so forth, all of which are made of solid matter, can be weighed, set on fire, demolished and towed away. Manhattan Island is not something that you can live *in* or walk *through*, though you are able to *fly over it*. The island and all its buildings, roads and so forth provide the needed material retainer, the physical support, for the many rooms in which New Yorkers can live and work, for the streets along which they can move, and for the tunnels through which they pass when riding the subway. Manhattan on this reading is a material

¹⁰ http://www.columbia.edu/cu/lweb/digital/collections/cul/texts/ldpd_6864674_000/index.html.

entity that is, like a Swiss cheese, full of holes. It has a complexity complementary to that of Manhattan under sense **1b**.

6. Making spatial conflicts

Prototypical fiat spatial entities are such that the results of the corresponding demarcations or parcelings out are accepted as a matter of course by the persons or organizations involved. Matters are not always so simple, however, as we know from the many examples of border conflicts in the course of history, from minor property-line disputes between neighbors to intercontinental warfare.

Consider Aksai Chin, a fiat spatial entity (compare **sense 1a** above) located where the borders of China and India meet. On maps printed in China, Aksai Chin is demarcated as belonging to China's Xinjiang and Tibet autonomous regions. On Indian maps, in contrast, Aksai Chin is demarcated as being included within India's borders. Both countries impose a legal requirement that Aksai Chin be represented correspondingly on all maps, and it is impossible for a third party (Burma, say) to produce a map that will not offend either China or India.

Under one scenario, one or other of the two parties to this dispute, let us say China, is correct in its assumptions related to the region Aksai Chin in virtue of the fact that, say, the United Nations has ruled in China's favour. China would be, under these circumstances, a *legal* fiat spatial entity. Even then, however, there would still exist a second fiat spatial entity, called 'India' by many Indians, which would be demarcated in such a way as to include Aksai Chin.

Under a second scenario, which is the scenario broadly accepted outside China and India today, neither the Chinese nor the Indian claim to Aksai Chin can be supported on objective grounds. Currently the area demarcated by the Line of Actual Control (LAC) coincides with the Chinese claim.

There is however no third scenario under which both China and India could simultaneously make a rightful and exclusive claim to the territory in question. The *legal* spatial entities now existing on the surface of the Earth are demarcated in such a way that, with very few exceptions, they form tessellations based on stable high-precision boundaries, with no overlaps¹¹, and no gaps¹².

7. Fiat spatiotemporal entities

¹¹ Pheasant Island/Île des Faisans/Isla de los Faisanes/Konpantzia, on the border between France and Spain, is administered jointly by the two countries, each of which is responsible for gardening on the island for 6 months in each year.

¹² There is no legally binding agreement as to where the borders between Switzerland, Austria, and Germany lie in the region of Lake Constance. Switzerland holds that the border runs through the middle of the lake; Austria that the lake is a condominium that belongs to all the states on its banks; Germany, meanwhile, holds the view that Germany holds no view on the matter.

Fiat spatiotemporal entities are one means by which we objectify past and future paths through space in order to reason about them. We all use such entities early on in planning a trip, as we engage in anticipatory manipulations of our trip plan, balancing in our minds the costs and benefits of alternative routes.

When we deploy online mapping software in the late stages of planning a journey, then we create a fiat spatiotemporal entity called a *route plan*, which is something like an abstract template that is then realized concretely in the journey we take. Because the software allows instantaneous manipulation of this plan, consequences of adjustments can be identified immediately. Once the plan has been fixed and agreed upon by those involved, however, then the route it dictates acquires something of the status of a *legal* fiat spatiotemporal entity projecting into the future. Commitments might be entered into on its behalf, for example in the form of purchases of theatre tickets or reservations at the campground, the latter marking out fiat spatiotemporal entities, namely *your seat in the theatre* or *your space at the campground* as these exist *during certain predefined periods of time*. This provides not just a way of allowing our rights to be exercised by ensuring that our reserved spaces will be available for us to use when we get there, but also a way of automatically coordinating the plans of multiple unconnected people by ensuring that their reserved spaces will be available to them when they arrive (Buckland 2008).

The demarcation of fiat spatiotemporal entities allows us also to manipulate experimentally regions of spacetime in the past. When historians seek to work out the route taken by Alexander the Great to reach the Indus River, then they create multiple drafts of fiat spatiotemporal entities that are projected onto the past as initial approximations on the one hand to the plans that Alexander may have entertained at different stages of his campaign, and on the other hand to the actual trajectory taken by his army. Something similar happens when population geneticists attempt to plot the evolutionary trajectories of ancient species; or when surgeons attempt to plot alternative minimally invasive surgical routes to a brain tumor in order to determine which would be the shortest. Fiat spatiotemporal entities are used in this manner also in attempts to understand the causes of a traffic accident, or of a security leak, or to work out the origins of a viral pandemic.

Conflicts can arise also between competing fiat spatiotemporal entities, for example when the alibi provided by a criminal suspect conflicts with accounts provided by witnesses. In 1977-1980, Soviet officials pressured Romania to edit the historical record to remove any references to the fact that most of the territory then comprising the Moldavian Soviet Socialist Republic had once been part of Romania. Here (at least) two conflicting fiat spatiotemporal entities are competing with each other for recognition as matters of historical fact.

8. Situated skilled action

There is an association between spatial entities and certain characteristic kinds of activities and concomitant human capabilities. This association applies, for example to the different zones in an airport (ticketing, boarding, transit, customs, ...), or to the different functional spaces in a courtroom (attorney tables, judge's bench, jury box, witness stand, ...). In 1950 the ecological psychologist Roger Barker coined the term 'behavior setting' to refer to spaces within which such characteristic activities occur (Wright and Barker 1950, Heft 2001, Smith 2000).

As Heft (2018) points out, when individuals join a behavior setting as participants, they

must have sufficient "know-how" for functioning in those settings, with "know-how" best characterized as *situated skilled action*. ... Because behavior settings arise from coordinated social practices, they have a normative character. Actions in a behavior setting must be appropriate in relation to the social practices that maintain it. (op. cit., p. 114)

An illustration of the way in which place and behaviour interact in this way is provided by those cases where football is played according to the rules on a field expressly demarcated for this purpose, with lines of paint marking out the zones of play and where persons other than players and officials are not allowed on the field while play is in progress.

When random groups of children are playing football in the schoolyard, in contrast, they may rely exclusively on the demarcation provided by the fact that they are using piles of coats as substitute goal posts. The boundaries of the field are otherwise to a large degree indeterminate; other children may occasionally cross the field of play, some of them merely taking a convenient shortcut to another point in the playground, others deliberately attempting to obstruct the course of play. Groups of children elsewhere in the schoolyard are engaging in multiple other sorts of activities. Different zones are hereby establishing themselves more or less spontaneously and then once more fading away in reflection of multiple forces, including the affinity and power relations between and within the different groups, but also including the stubborn fact that some of them are playing football and so it is expected that others will not impinge too heavily on their assumed space. Here and there there will be fights. When the schoolyard monitor appears some of the fights will immediately stop.

We can assume that the football game was able to be established and maintained because there are individuals involved who have been bequeathed by their developmental histories with strong desires to play football and also with developed capabilities to initiate and control and defend group activities of this sort. Other leader figures will have established similar islands of more or less coherent activity in other zones around the schoolyard. In engaging in these activities, and in dealing with the

occasional competitions for available space, the children are honing and testing out their capabilities to manage space, to track moving objects, to avoid crashing into other objects, to exert authority, and many more.

They are also honing and testing out their linguistic capabilities, for example the capability to shout and scream at invaders on the football field in a way that will minimize the disruption of play. They are thereby also training their brains, so that they will be better able in the future to address multiple simultaneous challenges of daily life in a maximally effective way.

9. The visual field and its horizon

Every organism with a perceptual system is creating a fiat spatial entity in every waking moment in the form of its *visual field* which reflects, very roughly, a demarcation between what it sees and what it does not see at any given time.¹³

There is no sharp boundary separating what is within from what is without, for the visual field is marked at its boundaries by a fading off between interior and exterior that is difficult to specify, and also by a continuous churning of what is contained in the field from moment to moment as we move our eyes and head or change the position or orientation of our whole body. We should therefore more properly talk of a succession of fields, determined by what we shall recognize below as the constant sensorimotor activity of perception. This succession is then analogous to the succession of snapshots created in the aerial survey referred to above. It is a succession of a type that is determined by the movements not of some target of perception, but rather of the perceiving/sensing subject.

This indeterminacy of borders involves also the phenomenon of what Husserl called the 'horizon'¹⁴ of the visual field (1989, pp. 142, 149). This is a matter not of what we see, but rather of the penumbra of possibilities, relating to what we might see, or do, or to what we expect (or want) to happen next, thus pointing ever forward.

The horizon is something like the penumbra of possible (next) perceptions. When we stand in front of a building it is not simply a plane surface (the front wall of the building) that we see. Rather, we see a *building*, a solid material object that extends in three dimensions. This is so because the sides and rear of the building are part of the horizon of our visual field because our experience incorporates the awareness that we can walk around the building and inspect it from any of a number of different

¹³ We focus here for simplicity on the visual field. In a more adequate accounting we would consider the spacemaking role of the entire sensory field, including what we can hear, and also what we can experience in proprioception.

¹⁴ A term we borrow from Husserl (1989, pp. 142,149).

perspectives.

10. Spacemaking in language

Spacemaking in language makes itself felt in our everyday lives when we follow directive signs ('No Smoking', 'No Children') or use expressions such as 'wait here!' or 'you can put your coat there', or 'we'll sit on this side of the table', or (perhaps accompanied by the tracing of one's finger across a map) 'we'll steer the boat down this channel' or 'we'll take this hill by nightfall!'. Such expressions have varying degrees of normativity, depending on the authority relations involved or assumed, and on the intentions of the speaker.

Fiat spatial entities are created also in the course of every conversation. First, conversations occur always in a certain spatial context, which is determined by the *site* or *setting* in which the conversation takes place, which may be a park bench, or a dining table, or a saloon bar.

The spatial setting of the conversation, too, will be marked in every case by its horizon, for example consisting of expectations concerning what our conversation partners might say next. We engage in conversation typically in order to realize certain intentions, and so our participation will be marked by the fact that we are aiming to achieve something by talking to these (or a certain subset of these) conversation partners. The anticipated incremental fulfilment or frustration of this aim will therefore also form part of our conversation horizon at any given stage.

There is now a two-way causal interaction between conversation setting and visual field, as what each conversation participant sees may influence the course of the conversation which may in turn influence what the other participants see, and so on down the line. What I say in a conversation and how say it will be influenced also by how the other participants are responding as I speak. This is because speech is a sensorimotor activity, and like other motor activities it is marked by the fact that when we perform an action we simultaneously obtain proprioceptive feedback from this very performance. This feedback will derive on the one hand from our body itself, and on the other hand from our environment via our sensory system. In the case of spoken language as used in dialogue experience first the feedback that derives from our hearing our own words as we are speaking. But this first type of feedback is then augmented by the feedback we receive from our interlocutor, for example in the form of facial expressions, gestures, as well as further speech. We continuously use this feedback to adjust not just what we say and the way we speak but also the intentions we are seeking to realize by engaging in dialogue.

11. How language structures space

In a series of classic papers¹⁵, Len Talmy has described the ways in which everyday language structures our cognition of space, and he identifies what he calls the spatial system of language, which

imposes a fixed form of structure on virtually every spatial scene. A scene cannot be represented directly at the one-structural level in just any way one might wish – say, as a complex of many components bearing a particular network of relations to each other. Rather, with its closed-class elements¹⁶ and the very structure of sentences, the system of language is to mark out one portion within a scene for primary focus and to characterize its spatial disposition in terms of a second portion ..., selected from the remainder of the scene. The primary object's spatial disposition here refers to its site when stationary, its path when moving, and often also its orientation during either state. (Talmy 2000, p. 181)

Here the 'scene' is what we have been referring to thus far as a fiat spatial entity, the 'path' is a fiat spatiotemporal entity. When Talmy describes how language structures space he is addressing, in the terminology we are using here, how almost every use of language creates fiat spatial and spatiotemporal entities. The difference is that we are focusing here primarily on how language structures (carves or demarcates portions of) *objective reality* as something that is spatially or spatiotemporally extended. Talmy, in contrast, is interested in the space-structuring abilities of language independently of any connection to the real world – thus his account applies equally well to the space-structuring features of language involved in works of fiction or in metaphorical uses of language as when I say 'he enveloped her with his love' or 'this mind was running in all directions'.

Talmy draws a tripartite division of a scene into Figure, Ground, and (sometimes) Secondary Reference Entity, which holds both where the Figure is stationary in space and where it moves through space. It is where a Figure moves that a path exists. A path is the cumulative set of locations that a moving Figure occupies during its motion, and its direction is determined by the sequence in which successive locations are occupied.

In his (1988) Talmy identifies a 'semantic category that has previously been neglected in linguistic study', namely that of force dynamics, which relates to

how entities interact with respect to force. Included here is the exertion of force, resistance to such a force, the overcoming of such a resistance, blockage of the expression of force, removal of such blockage, and the like. (Talmy 1988, p. 49)

¹⁵ These papers are collected in Talmy 2000, and supplemented by Talmy 2005)

¹⁶ Classes of words may be either open or closed. An open class is one that can easily accept the addition of new words. A closed class is one to which the addition of new words is very uncommon. Examples of such closed classes are pronouns, articles, conjunctions, and prepositions.

Here again, Talmy's focus is on the way in which force relations are expressed in language and are thereby understood cognitively or conceptually by the subjects involved. The phenomena of force dynamics thereby extend to the entire underlying human cognitive system. In our remarks on legal authority and related notions in the foregoing we have provided only a hint of some ingredients of an ontology of such phenomena – which in the realm of objective spatial reality would include possession, ownership and control, trespass, and ultimately invasion, defense, war, and so forth. Talmy's treatment of force dynamics is not concerned with applications of real physical force or with the threat of such applications that go hand in hand, for example, with the exertion of physical violence. His treatment will nonetheless provide an essential component in the extension of spatial ontology to these areas.

12. The limits of Artificial Intelligence

Our remarks on fiat spatial and spatiotemporal entities in the foregoing are intended as a contribution to the large literature on formal logical systems in the field of what is called qualitative spatial reasoning.¹⁷ Where this literature has concentrated overwhelmingly on issues relating to geographical space, we focus here on the idea of a general approach which can be applied also across legal/political, perceptual and linguistic domains, to planning as well as to history, as also to areas such as criminal evidence analysis. One component of this more general approach is the idea that spatial demarcations go hand in hand with special sorts of behaviour.

There are many sorts of conversations which occur in the sorts of *behaviour setting* we discussed above – behaviour settings in the sense of Roger Barker. These are conversations where all the participants share a similar set of capabilities and a substantial body of shared knowledge and shared presuppositions, for example where scientists are discussing results of recent experiments or politicians are discussing results of recent polls. Much of the publicly available data we have in the form of videos of debates or transcripts of interviews relate to conversations of this sort, in which the participants in addition observe standard rules of turn-taking, clarity of diction, and so forth.

Many conversations, however, evolve with no such shared setting because their participants have different social and cultural backgrounds, different education levels, are from different age cohorts, speak different dialects, or have a different attitude to rules of behaviour. Conversations of this sort may involve odd mixtures of background presuppositions; lead to misunderstandings; and will often fail because one or more of the participants will realize early on that their intentions will not be realized through further participation. Other features will add yet further variability, for example, hostile behaviour by one or more participants, sudden addition or departure of participants in mid-

¹⁷ For examples see (Galton 2000), (Varzi 2007), (Bateman et al. 2010) and (Bennett and Cohn 2021).

conversation, occasions where multiple people are speaking (or perhaps shouting) at once, or where the conversation turns upon itself (becomes a conversation, for example, about whether this very conversation should be taking place at all) or where interlocutors suddenly reach back to utterances made much earlier in the conversation and bring them once more into the foreground. (Such discontinuities in the conversation are from the mathematical point of view erratic (non-Markov).)

This huge array of different ways in which conversations may be structured – documented at greater length in (Landgrebe and Smith 2019) has implications for attempts to create AI systems that would be able to acquire the sort of language capability that humans manifest in their conversational behaviour. The currently popular approach to solving this problem involves the application of the technology of what are called ‘deep neural networks’, a technology whose tremendous power has been demonstrated most recently in areas such as playing GO and protein folding prediction. The technology works by exploiting features of stochastic processes which allow input-output patterns recognizable in huge samples of data to be used to ‘train’ a neural network in such a way that the network can recognize these same patterns across even larger spaces of inputs and thereby predict the corresponding outputs in useful ways. (This is how, for example, spam filters work.)

Unfortunately, due to the non-ergodic and evolutionary character of the human conversation system¹⁸, it is never possible to collect data from conversations that can be used to train a machine so well that it could become an acceptable conversation partner. This is because it is impossible to obtain the *representative data* which would be needed in order to train a neural network to make the necessary predictions. Due to the evolutionary character of the human mind and its ability to create new patterns at any moment, each conversation may evolve in such a way that it is *sui generis*.

Recall our discussion of children testing out their capabilities in the midst of schoolyard fights. When children react to threats by shouting and screaming, they do not do this after reasoning over available options or after any sort of conscious or unconscious computing or calculating. Rather, they act *spontaneously*, in a natural reaction that is analogous to the way in which flood water flows naturally across uneven terrain and never stops moving until the water has found its new level. In the brain, what is flowing are molecules with electrochemical properties and the uneven terrain across which they flow is a gigantic interwoven mesh of billions of molecules that are themselves constantly in flux as ever new experiences lay down new connections between them and introduce new and dispersed molecules into the flow. Language understanding in general can be seen as analogous to an automatic process of the sort whereby water flowing into uneven terrain will always, automatically, find its own

¹⁸ Or, more generally, of the human brain or central nervous system. In this context, ‘non-ergodic’ means that the elements of the conversation are never distributed according to a predictable pattern. A system is ‘evolutionary’ if new elements, new element types and new interactions and interaction types can arise at any time.

level. This capability for spontaneous understanding and of spontaneous responsive language production is of vital importance and reveals the natural, biological nature of especially spoken language as a sensorimotor phenomenon. But we understand almost nothing about how it works.

There is however an older, non-stochastic, logic-based approach to artificial intelligence, sometimes referred to as GOFAI (for 'good old-fashioned AI'), which was the dominant paradigm of AI research from the mid-1950s until the late 1980s (Haugeland 1985). In their *Rebooting AI*, Marcus and Davis (2019) outline an approach incorporating elements of GOFAI with the goal of supporting the creation of the sort of general artificial intelligence (AGI) that would be needed for mastery of human language on the part of the machine. The idea is to create a counterpart inside the machine of the sort of background or 'common sense' knowledge that humans bring as a matter of course to every conversation – and indeed to every interaction with other human beings and with their environments.

Humans have an impressive capability to resolve ambiguities in the course of conversation. When someone says *He put the ketchup on himself* then we know (instinctively, as Marcus and Davis would say) that this is about putting ketchup on food.

Nobody has to teach us how to do this sort of inference in school, because we know how to do it instinctively; it follows naturally from how we interpret the world in the first place. Deep learning can't even begin to tackle this sort of problem. (Marcus and Davis, 2019, p. 93)

And when it comes to capturing such background knowledge

classical AI researchers (and philosophers) have made excellent progress, working out formal logical systems for representing situations and how they develop and change over time. (op. cit., p. 162)

Will an approach like that suggested by Marcus and Davis help to bring us the sort of AGI they envisage? Unfortunately not – again because of the inexhaustible variability of human language use. It is impossible to create the sort of formal theories, even of just the spatial aspects of human language use, that will enable us to capture the meanings even of everyday natural language expressions. For as Bennett and Cohn (2021) show, even for relatively simple combinations of the sorts of spatially extended entities we encounter in everyday experience, there are potentially infinite distinctions among spatial relations.

A further problem, already well known to those who have attempted to develop systems for common-sense reasoning is the problem of computational complexity. Indeed, as Bennett and Cohn point out, many of the intractable or undecidable problems studied in computational complexity theory are spatial in nature.

13. Conclusion

To what purpose, then, are we developing spatial ontologies based on obscure ideas of *fiat spatial ad spatiotemporal entities*. One answer is precisely the utility of suitably tailored ontologies in achieving at least a useful part of what protagonists of AGI are striving for.

In (Landgrebe and Smith 2021, Appendix) we describe a multi-stage multi-component strategy to use logically axiomatized formal ontologies, along with stochastic algorithms and other components to translate complex natural language expressions in a way that will capture in logical form the full meaning of these expressions as used in a given domain and for a given purpose.

A *general purpose* application of this sort is, by our arguments above, impossible. If, however, we focus on specific narrow domains, where there is a limited, and well-understood vocabulary, then we can build corresponding domain-specific ontologies that can be deployed along the lines described to generate the representation in machine-readable form not merely of single expressions but of entire texts, even of entire corpora of texts. The outputs of this process are logical formulae using ontology terms which are counterparts of the constituent simple expressions used in the relevant domain.

The methodology has been applied, for example, to the automatic processing of correspondence between customers and energy suppliers. Where the system cannot respond adequately to a given piece of correspondence, the item in question is passed on for handling by a human expert.

We believe that systems of this sort could be developed in areas such as insurance claims processing, legal text analysis, medical billing, analysis of descriptions of financial assets (such as bonds and securities), or in domain-specific semantic search.

We emphasize once more that the strategy can be used only in narrowly restricted domains, with a narrowly restricted vocabulary, and only after the manual construction, testing and refining of detailed ontologies. We create general ontologies, for example of spatial entities and spatial relations, not because we believe that such ontologies can support general AI, but because the difficult task of building high-quality detail-focused domain-specific ontologies is made considerably easier if we have a consistent set of mature, well-tested more general ontologies that can provide a reliable starting point for domain ontology development and potentially also increase the degree to which expertise derived from building ontologies in one domain can be applied also in neighbouring domains. This final consideration is important not least because many of the domain ontologies built in recent years failed because their developers did not draw on any widely accepted and well-understood principles of development.¹⁹

¹⁹ With thanks to Harry Heft, Len Talmy and Brendan Whyte for many valuable comments.

Compliance with Ethical Standards:

Author declares that he has no conflict of interest.

Ethical approval: This article does not contain any studies with human participants or animals performed by any of the authors.

References

Arp R, Smith B, Spear AD (2015) *Building ontologies with Basic Formal Ontology*. MIT Press, Cambridge MA.

Barker RG (1968) *Ecological psychology*. Stanford University Press.

Bateman JA., Hois J, Ross R, Tenbrink T (2010) A linguistic ontology of space for natural language processing. *Artificial Intelligence*, 174(14):1027-1071.

Bennett B (2010). Foundations for an ontology of environment and habitat. In *Formal Ontology in Information Systems (FOIS)*, pp. 31-44.

Bennett B, Cohn A (2021) Automated common-sense spatial reasoning: still a huge challenge. In: Muggleton S, Chater N (eds) *Human-like machine intelligence*, Oxford University Press.

Buckland M (2018) Document theory. *Knowledge Organization*, 45(5):425-436.

de Soto, H (2000) *The mystery of capital: Why capitalism triumphs in the West and fails everywhere else*. Civitas Books, New York.

Dreyfus HL, Dreyfus SE (2005) Peripheral vision: Expertise in real world contexts. *Organization studies*, 26(5):779-792. <https://doi.org/10.1177/0170840605053102>

Fónagy I (2001) *Languages within language. An evolutive approach*. John Benjamins, Amsterdam/Philadelphia.

Froehlich K (1981) The plan of St. Gall, *Theology Today*, 38(1):84-91.

Galton A (2000) *Qualitative spatial change*. Oxford University Press.

Gibson JJ (1966) *The senses considered as perceptual systems*. Allen and Unwin, London.

Haugeland, J (1985) *Artificial Intelligence: The very idea*. MIT Press, Cambridge, Mass.

Heft H (2018) Places: Widening the scope of an ecological approach to perception–action with an emphasis on child development, *Ecological Psychology*, 30:99-123.

<https://doi.org/10.1080/10407413.2018.1410045>

Heft H (2001) *Ecological psychology in context: James Gibson, Roger Barker, and the legacy of William James's radical empiricism*. Psychology Press, New York.

Husserl E (1989) Husserl, Edmund. 1989. *The crisis of European sciences and transcendental phenomenology: An introduction to phenomenological philosophy*. Northwestern University Press, Evanston, Illinois.

Landgrebe J, Smith B (2019). There is no Artificial General Intelligence. arXiv preprint. arXiv:1906.05833.

Landgrebe J, Smith B (2021) Making AI meaningful again. *Synthese*, 198:2061–2081.

<https://doi.org/10.1007/s11229-019-02192-y>

Marcus G, Davis E (2019) *Rebooting AI: Building artificial intelligence we can trust*. Vintage, New York.

Mark DM, Smith B (2003) Do mountains exist? Ontology of landforms and topography.

Environment & Planning B (Planning and Design), 30(3):411-427. <https://doi.org/10.1068/b12821>

Mark DM, Smith B (2004) A science of topography: From qualitative ontology to digital representations. In: Bishop MP and Shroder JF (eds) *Geographic Information Science and Mountain Geomorphology*, Springer-Praxis, Chichester, England, 2004:75-100.

Mulligan K (1987) Promisings and other social acts: Their constituents and structure. In K Mulligan (ed) *Speech act and Sachverhalt. Reinach and the foundations of realist phenomenology* (pp. 29-90). Springer, Dordrecht.

Reinach A (2013). *The a priori foundations of the civil law*. Translated by J Crosby. W De Gruyter, Frankfurt.

Searle JR (1964), How to derive 'ought' from 'is'. *Philosophical Review*, 73:43-58.

Searle JR (1969) *Speech acts: An essay in the philosophy of language*. Cambridge University Press.

Searle JR (1995) *The construction of social reality*. Simon and Schuster, New York.

Searle JR (2010). *Making the social world: The structure of human civilization*. Oxford University Press.

Smith B (1995) On drawing lines on a map. In: Frank AU and W Kuhn (eds), *Spatial information theory. A theoretical basis for GIS*, Springer, Berlin/Heidelberg/New York, pp. 475-484.

https://doi.org/10.1007/3-540-60392-1_31

Smith B (2000) Objects and their environments: from Aristotle to ecological ontology. In *Life and motion of socio-economic units*, pp. 84-102. CRC Press, Boca Raton.

Smith. B (2001) Fiat objects, *Topoi* 20(2):131-148. <https://doi.org/10.1023/A:1017948522031>

Smith B (2008) Searle and de Soto: The new ontology of the social world. In Smith B, Mark DM and Ehrlich I (eds) *The mystery of capital and the construction of social reality*, Open Court, Chicago, pp. 35-51.

Smith B, Searle JR (2003) The construction of social reality: An exchange. *American Journal of Economics and Sociology*, 62(2):285-309.

Smith B, Varzi AC (1999). The niche. *Noûs*, 33(2):214-238.

Smith B, Varzi AC (2002). Surrounding space – The ontology of organism-environment relations. *Theory in Biosciences*, 121(2):139-162.

Talmy L (1983) How language structures space, in: H. Pick, L. Acredolo (Eds.), *Spatial orientation: Theory, research, and application*, Plenum Press, New York, pp. 225–282.

Talmy L (1988) Force dynamics in language and cognition. *Cognitive science*, 12(1):49-100.

Talmy, L. (2000). *Toward a cognitive semantics*, vol. 1. Language, speech, and communication. MIT Press, Cambridge, MA.

Talmy L (2005) The fundamental system of spatial schemas in language. In B Hampe (ed), *From perception to meaning: image schemas in cognitive linguistics*, Mouton de Gruyter, Berlin, pp. 37–47. <https://doi.org/10.1515/9783110197532>

Varzi AC (2007) Spatial reasoning and ontology: Parts, wholes, and locations. In *Handbook of spatial logics* (pp. 945-1038). Springer, Dordrecht.

Wright HF, Barker RC (1950). *Methods in psychological ecology; a progress report*. Midwest Field Study of Children's Behavior.