

Qualitative spatial representation in agent-based models

Gary Polhill
The James Hutton Institute,
Craigiebuckler, Aberdeen,
AB15 8QH. UK
Email: gary.polhill@hutton.ac.uk

Nick Gotts
Independent researcher
Email: ngotts@gn.apc.org

Amparo Alonso-Betanzos,
Noelia Sanchez-Maroto,
Óscar Fontenla-Romero
Facultad de Informática,
Universidade da Coruña,
A Coruña 15071, Spain
Email: {[ciamparo](mailto:ciamparo@udc.es), [nsanchez](mailto:nsanchez@udc.es),
[ofontenla](mailto:ofontenla@udc.es)}

Tony Craig
The James Hutton Institute,
Craigiebuckler, Aberdeen,
AB15 8QH. UK
Email: tony.craig@hutton.ac.uk

Ricardo García-Mira
Departamento de Psicología,
Universidade da Coruña,
A Coruña 15071, Spain
Email: ricardo.garcia.mira@udc.es

Abstract—One of the advantages of agent-based models as simulations of social systems is the ease with which it is possible to spatially embed the agents and their interactions. Spatially explicit representations in agent-based models most typically take the form of raster-based representations in which the space is represented as a grid of squares. More recently, vector-based representations have been used, usually importing data for the polygons from geographical information systems (GIS). However, for some models, what matters about the space for the purposes of simulation is less the quantitative spatial relationships among entities (e.g. area, distance or direction) than the qualitative relations these quantitative data are used to determine: neighbourhood, and accessibility (which is a general term covering movement and sensing from one region to another). This paper gives consideration to the use of qualitative spatial representations in agent-based modelling, using a model of everyday pro-environmental behaviour in the workplace as an example.

I. INTRODUCTION

SPATIAL representation in agent-based models typically consists of raster-based representations, if only because popular tools and libraries such as NetLogo [1] and Repast [2] provide default utilities to do so. Neighbourhood functions use a shared boundary or point as a basis, leading respectively to the von Neumann neighbourhood (the four cells sharing a border with a given cell; `neighbors4` in NetLogo), or the Moore neighbourhood (the eight cells sharing a border or a corner with a given cell; `neighbors` in NetLogo) interpretations of cell-based neighbourhood. Vector-based representation, in which the space is divided into irregular polygons (typically, though not necessarily (see

[3]) corresponding to real geographical features), is increasingly being used [e.g. 4]. This is true particularly of models with a geographical application, as data may be conveniently read in from a GIS. Indeed, NetLogo provides the GIS extension, which reads ESRI shape files (a popularly used format for vector data), and Repast includes packages and classes to enable modellers to work with GIS data. A raster-based representation can be seen as a special case of a vector-based one in which the polygons are all squares of the same area, forming a grid.

Both of these representations are quantitative, in the sense that the vertices of the polygons are defined using numerical co-ordinates which, in Euclidean spaces, allows the computation of distances between pairs of points and directions from one point to another using the standard formulae of geometry. However, in agent-based models, it is often the relationships that these spatial entities define that are most important: where the agent can move, what the agent can sense, and what the agent can interact with. In some cases, quantitative spatial representations may arguably constitute spurious precision, and create a misleading impression of the accuracy with which the spatial aspects of the model are predicted.

Qualitative spatial representations have been the subject of research for a number of years, the Region Connection Calculus (RCC) [5] being a popular formalism. Qualitative representations capture logical relationships among regions or cells in a space, rather than relying on their quantitative aspects (i.e. bounding co-ordinates) to derive them.

In this paper, we consider the applicability of RCC to agent-based models. Reflecting on the requirements of agent-based models, we find that extensions of these formalisms will be needed if qualitative spatial representations are to have a role in ABMs in the future. We suggest some prototype extensions related to agents' sense perceptions and

[≡] This work was jointly funded by the European Commission through Framework Programme 7 grant agreement number 265155 (LOCAW: Low Carbon at Work: Modelling Agents and Organisations to Achieve Transition to a Low-Carbon Europe), and by the Scottish Government Rural Affairs and the Environment Portfolio Strategic Research Theme 4 (Economic Adaptation).

movement, and illustrate the use of one of them in a model of everyday pro-environmental behaviour at work.

II. RCC AND QUALITATIVE SPATIAL REQUIREMENTS FOR ABM

RCC builds relations among regions from the primitive mereotopological relation C (connected). Although RCC was originally conceived of as a 'pointless' representation of space [5], regions can be thought of as non-empty open subsets of a universal region conceived of as a set of points [6]. C can be used to define P (part) by saying that region x is part of region y if all regions connected to x are connected to y :

$$P(x, y) \equiv \forall z (C(z, x) \rightarrow C(z, y))$$

P can then be used to define O (overlap) between x and y as the existence of a region z that is part of both x and y :

$$O(x, y) \equiv \exists z (P(z, x) \wedge P(z, y))$$

There are then a number of different ways of constructing sets of jointly exhaustive pairwise disjoint relations among regions. RCC-5 defines the five relations DR (disjoint regions), PO (partially overlaps), PP (proper part), PPi (inverse proper part) and EQ (equal) as:

$$DR(x, y) \equiv \neg O(x, y)$$

$$PO(x, y) \equiv O(x, y) \wedge \neg P(x, y) \wedge \neg P(y, x)$$

$$PP(x, y) \equiv P(x, y) \wedge \neg P(y, x)$$

$$PPi(x, y) \equiv PP(y, x)$$

$$EQ(x, y) \equiv P(x, y) \wedge P(y, x)$$

Note that PO represents what might be more commonly understood by the term 'overlaps' in that (unlike O) it does not apply when one of the related regions is wholly contained within the other. Similarly, PP (as in standard mereology) represents the idea that might normally be understood as 'part', in that it precludes equality of its arguments. (From the above definitions, it is easy to see that PP could equivalently be defined as " P and not EQ ".)

RCC-8 further discriminates DR into DC (disconnected) and EC (externally connected), and PP into TPP (tangential proper part) and $NTPP$ (non-tangential proper part), with inverses $TPPi$ and $NTPPi$:

$$DC(x, y) \equiv \neg C(x, y)$$

$$EC(x, y) \equiv C(x, y) \wedge \neg O(x, y)$$

$$TPP(x, y) \equiv PP(x, y) \wedge \exists z (EC(x, z) \wedge EC(y, z))$$

$$NTPP(x, y) \equiv PP(x, y) \wedge \neg \exists z (EC(x, z) \wedge EC(y, z))$$

Although regions in RCC have no concept of a boundary, EC can be conceived of as the closures of x and y intersecting whilst the open sets x and y constituting the regions themselves do not intersect.

RCC, because of the purposes for which it was designed, focuses exclusively on topological relationships among spatial regions. Embedding agents in space inevitably creates the need for new relationships, which take two forms: one defining relationships agents have with places (e.g. ownership, location); the other defining how relationships among agents are mediated through and interact with space. We consider the former here.

Links between description logics and RCC have already been established through attempts to represent it in OWL [7, 8]. Since there are interactions between the ontology of the model and that of its (qualitative) space, this holds out the hope that existing formalisms could be exploited to achieve the expressiveness needed to meet requirements for qualitative spatial representation in ABMs, particularly since not all of the expressiveness of RCC is necessarily needed for this purpose.

As suggested in the introduction, there are three main areas where agent-based models use space, though specific applications may vary: neighbourhood, sensing and movement. The last two are closely related, and may be considered together under the more general heading of accessibility.

A. Neighbourhood

Neighbourhood in a topological sense is captured by the RCC-8 relation EC : disjoint regions that are connected but do not overlap. As noted in early work with FEARLUS [9], there can be a distinction between topological and social neighbourhood. Suppose a relation Q holds between an agent and a region (such as ownership). Then the social neighbourhood N with respect to Q could be defined in terms of Q and EC thus:

$$N(a, b) \equiv \exists x, y (Q(a, x) \wedge EC(x, y) \wedge Q(b, y) \wedge a \neq b)$$

Indeed, Q forms the basis of an ontologically significant spatial scale, in that a region can be defined as the sum of those regions for which Q holds for a particular agent (i.e., those that the agent owns for the purposes of this example). (RCC does not stipulate continuity of regions.) However, since it seems desirable (for relationships such as ownership) to assert that if Q holds for a region then it also holds for its parts:

$$Q(a, x) \wedge P(y, x) \rightarrow Q(a, y) \quad ,$$

a spatial scale should only apply to sets of regions that are not parts of regions with Q to the same agent. Specifically, we can define a spatial scale with respect to Q as the set of maximal regions S for which Q holds for some agent, where 'maximal' means that members of S are not proper parts of regions for which Q holds:

$$S = \{x : \exists a(Q(a, x) \wedge \neg \exists y(PP(x, y) \wedge Q(a, y)))\}$$

Although in general Q could be a many-to-many relationship, if Q is an inverse functional relationship (i.e. each region can be related by Q to just one agent, as should be the case with ownership), then we know that no member of S will overlap with any other member of S .

If there exists another relation R between agents and regions (a subproperty or specialisation of Q) such that:

$$\forall a, x(R(a, x) \rightarrow Q(a, x)) \quad ,$$

then we should expect that regions for which R holds are parts of regions for which Q holds. The hierarchy of ontological relationships connecting agents to regions in various different ways form a matching partial ordering of spatial scale.

B. Accessibility

One of the issues with movement for qualitative spatial reasoning, given the lack of any quantification of distance, is its relationship with time. Typically, agent-based models feature discrete time steps that are quantitative at least in the sense that they correspond at some level to real-world temporal intervals such as days, months or years. As such, there may be arguments that qualitative spatial representations should only be used with qualitative temporal representations (i.e. simply knowing the (partial) order in which events take place). The knowledge that nothing can travel faster than the speed of light, for example, means that if an agent moves from one region to another in a day, then the distance travelled must be less than 25.9 Tm. Though somewhat extreme, this does mean that the quantification of time allows a 'creep' of quantification into the representation of space.

That issue aside, the location of an agent is another relationship it may have with space. In contrast with the example of ownership explored above, if an agent is located in a region, we would also say that it is located in regions of which that region is a part, but we would not (necessarily) say that the agent is in parts of that region. For example, if Geert is in his office, he is not necessarily in all parts of that office, but he is on the floor on which the office is located, and in the building of which that floor is a part. Using L to represent the location of an agent, note the contrast with Q :

$$L(a, x) \wedge P(x, y) \rightarrow L(a, y)$$

It would be convenient if we could identify a set of regions in which agents could be located, as was the case for S above with respect to Q . A number of issues make this a challenge, not least of which is the fact that the set of regions in which agents are currently located does not exhaustively specify the set of regions in which they *could be* located. Further, the question of which regions agents can meaningfully locate themselves in is part of the narrative of the model, and hence the set of regions may be part of the configuration rather than something that is inferred from other ax-

ioms. Finally, it may be required to explicitly represent the *embodiment* of an agent as a region, and the location of the agent is then those regions of which that embodiment is a part.

For the purposes of movement from one region to another, it may be reasonable to stipulate minimally that the two regions are related by EC . Equally, it seems reasonable to argue that if an agent is located in a region that partially-overlaps with another region that the agent is not currently located in, then the agent can move between the two regions. Such arguments do not allow for the possibility of regions acting as obstacles for movement, however, and there is the further question of whether there is heterogeneity in agents' capability to move between regions. In general, it may be simpler to identify as part of the model specification a set of regions in which agents may be described as being located, and then to specify a relation identifying those pairs of regions between which an agent can move in a single time step. Where there are differences in agents' capabilities for movement, this relation will have to be specialised according to the class of agent (where this class is defined using appropriate combinations of restrictions on the agents' attributes).


Sensing is closely related to movement in terms of the issues raised, in particular obstacles to sensing, and differences in agents' capabilities to sense. There is also a complicated relationship with time with respect to the synchrony of events in that knowing that one agent can sense another depends on knowing where they are and when, bearing in mind that for some (stigmergic) senses, it is not necessarily the case that the object of the sensing is present within range at the time the subject performs the detection. Stigmergic sensation (such as ant trails, or other chemical or physical alterations of the space) is an interaction among agents that is mediated through space, and requires that regions have (physical) attributes that the agents can detect and modify, a matter that RCC does not address.

Two physical attributes are of potential use in discussing the obstacles to sensing: opacity and transparency. Asserting the transparency of a region could be a way to stipulate that it offers no obstacle to sensing from one region to another of which it is a part. Similarly, stating that a region is opaque could be used to assert that no sensing is possible between agents located in that region. The latter could be used for regions that are the sum of disjoint regions, such as the region comprised of all the houses of employees of a company.


III. PARTIAL DEMONSTRATION IN WERC-M

WERC-M (Worker-Environment Reinforcement Choice Model) is a model of everyday pro-environmental behaviour at work created for the LOCAW Framework Programme 7 project.¹ It was built to model backcasting scenarios aimed at improving everyday pro-environmental behaviour in four case studies of workplaces in the utility and public sector (Aquatim, a water utility in Romania, ENEL Green Power, an electricity company in Italy, Groningen municipality in the Netherlands and the University of A Coruña in Spain).

¹ Low Carbon at Work: Modelling Agents and Organisations to achieve Transition to a Low-Carbon Europe. <http://www.locaw-fp7.com/>



ig 1: Regions shaded to show they apply to switching lights off in one of the office buildings



ig 2: Regions applying to the using paper cups context




Fig 3: Regions applying to short business trips context




Fig 4: Regions applying to short business trips context after implementing a home-working scenario

A 'context' in WERC-M is a situation in which the agent has to make a decision about whether to behave pro-environmentally or not. Geller et al. [10] have a specification of contexts that is relevant to the use of the term here. This provides details of such things as the conditions under which a context is initialised (which can include another context, as is the case here), the actors affected, and decision making: the actions the agents can perform in the context (here: behave pro-environmentally or not), and how the agent choses which action it will take (here, the appropriate decision tree). However, the specification does not include spatial informa-

tion, and since some of the contexts in WERC-M apply in specific region types (e.g. meeting rooms), we recommend an update to the context specification that incorporates these considerations. However, the distinction between contexts (situations in which agents have to make a decision) and regions (areas of space in which agents are located and contexts may occur) should be clear.

Detailed quantitative geographies of the workplace layouts were not available from the case study partners, and the model relied on knowing who could see whom for transmission of injunctive and descriptive norms [11]. Descriptive

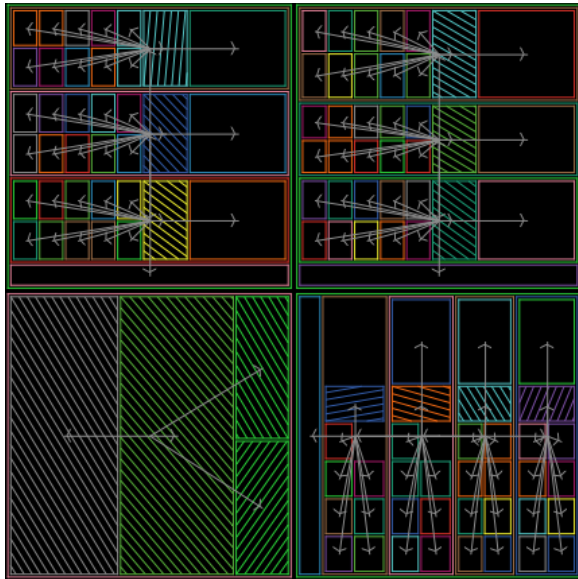


Fig 5: Depiction of the regions and part links in WERC-M. Arrows point from the centre of regions to their parts. Shaded regions have `opaque? = true`.

norms are related to imitation, and rely on the agent being able to observe others to find out what they do. Injunctive norms involve the agent being told what to do, and rely on the agent being seen behaving in such a way that an instruction should be given.

We therefore gave consideration to the use of qualitative spatial representations in the model, defining workplace environments such as the canteen, offices and open-plan areas, and, since the model also had to consider the possibility of spillover effects, home regions. Two qualitative spatial relations proved important: the mereotopological relation *proper part* (*PP*) in defining where contexts occurred that required a decision to be made about whether to behave pro-environmentally, and a sight relation determining who could observe whom (whether for injunctive or descriptive norm purposes).

The proper part relation allows contexts to be associated with places and all places that are parts of those places. Certain such contexts are associated with the agents' workplace (office or desk), such as switching off the computer at the end of the day, using emails rather than paper-based communication, or adjusting heating or lighting. Other contexts are associated with canteens, such as deciding whether to use a paper cup for drinks; or with home spaces, such as decisions about whether to wash clothes with a full load, or to dry them outside.

In the NetLogo implementation of WERC-M, we defined regions as agents (in the 'turtle' sense of the word), and created the directed link breed `part` to represent *PP*. In Fig 5, the regions for the Netherlands case study are shown, each as a rectangle with a distinct border colour. Regions contain their parts in the diagram. The home region is in the bottom-left quarter of the diagram, divided into four subregions: homes with and without air conditioning on the left and in middle, and on the right hand side, home offices with

and without air conditioning on the top and bottom. The three remaining quarters of the diagram each correspond to three separate office buildings, each of which is divided further into a number of floors, with each floor containing an open plan area, or a number of shared and unshared offices and a kitchen for preparing drinks.

In Fig 1, the `part` link is used to show how an everyday behavioural context (switching off the lights) applying to a specific building applies to all regions that are parts of that building. By tying the context to that particular building (rather than the workplace in general), we can implement scenarios where automatic lighting, which removes the need for an agent to make a decision to switch the lights off manually, is installed in different buildings at different times. In the diagram, regions are shaded using a number of parallel lines with a random angle; hence cross-hatching indicates subregions. In the NetLogo model, we also assigned regions a `region-type`, to discriminate between home regions, work regions, and 'third' spaces (areas in the workplace where norms associated with leisure or home might be more likely to apply). Kitchens are examples of such 'third' spaces, where interactions among colleagues may be more informal, or pertain to less work-related matters than office or open-plan areas within the building. Fig 2 shows the kitchen areas, shaded because they have been determined automatically to apply to the context associated with deciding whether to use your own cup or a disposable paper cup in the coffee machine.

Agents in WERC-M have an assigned `work-region` indicating their designated place of work, which is a potentially many-many relationship akin to Q discussed earlier. Contexts in WERC-M can be designated as applying to agents in their `work-region`, allowing intervention scenarios aimed at increasing pro-environmental behaviour to explore changing the `work-region`. An example of this is increasing the use of home-working. Fig 3 and Fig 4 show respectively the regions affected by the context of choosing a mode of transport for a short business trip (of less than five kilometres) before and after implementing a policy in which 10% of the workforce normally work from home.

Movement in WERC-M is represented simply using routines, in which the day is divided up into a series of time chunks (not necessarily of equal length), each of which is associated with a `region` or `region-type` in which agents fulfilling particular roles in the organisation will be located. Movement between certain pairs of regions (such as those between a work region and a home region) as the model proceeds from one time chunk to the next can then act as triggers for contexts to occur (such as commuting).

To address the question of sight, the model gave regions the Boolean properties `transparent?` and `opaque?`. The `transparent?` property was intended for the purposes of providing for open-plan areas, in which agents would have a designated region for their workplace (a desk), but would be able to see other agents in the open-plan area. Here, a desk would be defined as a `transparent?` subregion of an open-plan area, and an agent at that desk would be able to see agents in all `transparent?` subregions (i.e.

other desks) of the open-plan area. The principle is shown in Fig 6, where shared offices have been assigned transparent? = true (for the purposes of illustration), and sight links are created between pairs of regions for which agents in one region can observe agents in the other.

However, in the end it was simpler just to say that everyone worked in the open-plan area, without being specific about where. Of more use was the opaque? property, which is used to prevent agents seeing each other if they are in the same region. This allowed the home to be represented as a single region for all agents to live in, whilst ensuring no agent could observe another therein. Similarly, a single region could be used to represent all single-occupancy offices on a floor – the opaque? = true setting preventing those occupants observing each others' behaviour while located there. (See Fig 5.)

IV. CONCLUSION

Qualitative spatial reasoning has a potentially important role to play in agent-based models, but work is needed to define suitable formalisms for agent-based models to use so that they can be included in popularly-used agent-based modelling tools and libraries. We have shown how some of these principles can be implemented in NetLogo using existing functionality. In extending existing formalisms for qualitative spatial reasoning, there is the opportunity to draw on social theories defining the relationships between humans and the space they inhabit, and the constraints space imposes on relationships among humans.

REFERENCES

- [1] U. Wilenski, *NetLogo*. Center for Connected Learning and Computer-Based Modeling, 1999. <http://ccl.northwestern.edu/netlogo/>
- [2] M. J. North, N. T. Collier and J. R. Vos, "Experiences creating three implementations of the Repast agent modeling toolkit," *ACM Transactions on Modeling and Computer Simulation*, vol. 16, iss. 1, pp. 1-25, 2006.
- [3] A. Flache and R. Hegselmann, "Do irregular grids make a difference? Relaxing the spatial regularity assumption in cellular models of social dynamics," *Journal of Artificial Societies and Social Simulation*, vol. 4, iss. 4, p. 6, 2001. <http://jasss.soc.surrey.ac.uk/4/4/6.html>
- [4] A. T. Crooks, "Constructing and implementing an agent-based model of residential segregation through vector GIS," *International Journal of Geographical Information Science*, vol. 24, pp. 661-675, 2010.
- [5] A. G. Cohn, B. Bennett, J. Gooday and N. M. Gotts, "Qualitative spatial representation and reasoning with the Region Connection Calculus," *Geoinformatica*, vol. 1, pp. 1-44, 1997.
- [6] N. M. Gotts, "An axiomatic approach to topology for spatial information systems," *Research Report 96.25, School of Computer Studies, University of Leeds*, 1996 https://www.engineering.leeds.ac.uk/computing/research/publications/reports/1996/1996_25.pdf
- [7] R. Grütter and B. Bauer-Messmer, "Combining OWL with RCC for spatioterminal reasoning on environmental data," in C. Golbreich, A. Kalyanpur and B. Parsia (eds.) *Proceedings of the OWLED 2007 Workshop on OWL: Experiences and Directions, Innsbruck, Austria, June 6-7, 2007*. <http://ceur-ws.org/Vol-258/paper17.pdf>
- [8] R. Grütter, T. Scharrenbach and B. Bauer-Messmer, "Improving an RCC-derived geospatial approximation by OWL axioms," in A. Sheth, S. Staab, M. Dean, M. Paolucci, D. Maynard, T. Finin et al. (eds.) *The Semantic Web -- ISWC 2008, 7th International Semantic Web Conference, ISWC 2008, Karlsruhe, Germany, October 26-30, 2008, Proceedings. Lecture Notes in Computer Science* vol. 5318, pp. 293-306.
- [9] J. G. Polhill, N. M. Gotts and A. N. R. Law, "Imitative versus nonimitative strategies in a land-use simulation," *Cybernetics and Systems* vol. 32, pp. 285-307.
- [10] A. Geller, M. M. Latek, S. M. Mussavi Rizi, S. Fournier and F. Prenot-Guinard (2012) Contexts as reasoning and action frames for multiagent societies. <http://computationalsocialscience.org/wp-content/uploads/2012/09/GellerEtAl2012.pdf>
- [11] R. B. Cialdini, R. R. Reno and C. A. Kallgren, "A focus theory of normative conduct: Recycling the concept of norms to reduce littering in public places," *Journal of Personality and Social Psychology* vol. 58, pp. 1015-1026.




Fig 6: Depiction of sight links when shared offices are transparent?