Qualitative Reasoning with Spatially Distributed Parameters

Monika Lundell *
Artificial Intelligence Laboratory
Computer Science Department
Swiss Federal Institute of Technology
IN–Ecublens, 1015 Lausanne, Switzerland
E-mail: lundell@lia.di.epfl.ch

Abstract

This paper presents work on modelling the qualitative behaviour of physical systems of spatially distributed parameters. The distribution of each parameter is given as a set of observation points. A metric diagram is constructed by defining a connectivity structure on the point set. The metric diagram is used to construct a topological map that represents the distribution of the parameter as a pattern of contiguous regions. The regions represent values in the parameter's quantity space, which is a discretization of its value domain. Topological combinations of parameter distributions are used to infer the distributions of non-observed parameters according to models of parameter correspondences, e.g. qualitative versions of equations. The spatial evolution of the system is inferred by matching the scenario's parameter patterns against modelled patterns of physical processes. The approach is suitable for modelling common-sense reasoning in the natural sciences, e.g. meteorology, agriculture, climate studies and natural resource management.

1 Introduction

Qualitative models of physical systems often focus on describing how various parameters will evolve in *time*. The work described in this paper is also concerned with how parameters evolve in *space*. We present work on modelling the qualitative behaviour of physical systems of interacting spatially distributed parameters.

A *distributed* parameter is one that takes on different values at different points in space as well as in time. Many parameters describing the physical world can be modelled as distributed, e.g. temperature, colour, vegetation type, etc.

The human part of modern weather prediction is a good example of the kind of common-sense reasoning about spatially distributed physical systems we want to model. The role of the meteorologist is to analyze, understand and, if possible, predict the behaviour of the spatially distributed parameters of the atmosphere. The tools are a mixture of quantitative and qualitative methods. We will briefly outline some important steps:

- **Data collection:** Some key parameters in the atmosphere, e.g. temperature, air pressure and rain fall, are regularly and simultaneously measured at a number of observing stations.
- **Objective analysis:** The collected data is fed into a central computer where a numerical model is used to calculate a prediction for a large geographical region, in general for the next 24–72 hours. The numerical model uses some key physical laws in the form of differential equations, but contains many simplifications in order to make it tractable.
- **Subjective analysis:** The collected and predicted data is plotted on separate weather maps that are analyzed by hand by the meteorologist.
- **Prediction:** The meteorologist makes a prediction based on both the subjective and objective analyses. Most predictions concern short time periods and limited geographical regions that are not specifically catered for in the objective analysis, e.g. the area around an airport in the next hour.

From this description, we see that the computer-supported number-crunching of the objective analysis is only one part of the weather prediction process. We are interested in modelling the common-sense reasoning that takes place in the last two phases, i.e. subjective analysis and prediction.

Good weather predictions are based on a thorough understanding of the on-going physical processes in

*This research is supported by the Swiss National Science Foundation, project no. 5003-034269.

the atmosphere. The subjective analysis is time-consuming but necessary in order to build a mental model of these processes. This model is called "the inner weather picture" in [Perby, 1988], where the mental processes underlying weather prediction are discussed in more detail.

The subjective analysis starts with a study of the spatial distribution of each observed parameter. The observed values are indicated as points on a weather map. The meteorologist analyzes one parameter at a time by indicating regions of similar values on the map, e.g. isobars, isotherms, regions of precipitation, fog, cloudiness, etc. The analyzed map is used as a means of communication between meteorologists, and enables them to detect significant patterns of regions that indicate which underlying physical processes are at work. This process-based understanding creates an expectation of how the situation will develop, which is compared with the prediction of the objective analysis. The final prediction is based on the meteorologist's total understanding of the situation, which has been created from various sources: knowledge of physics, previous experience, collected data, objective prediction.

This kind of reasoning is interesting to artificial intelligence research as it involves at least four different research areas:

- **Model-based reasoning:** The reasoning is based on underlying models of physical phenomena.
- **Spatial reasoning:** The location where a parameter is observed is as important as the measured value.
- **Qualitative reasoning:** Due to the sparseness of observed data, assumptions and simplifications of a qualitative nature are necessary.
- **Diagrammatic reasoning:** Diagrams are extensively used to understand complex situations and to communicate this understanding.

We believe that the working methods of meteorologists are representative of many other scientific areas where physical systems of spatially distributed parameters are studied. Some examples are natural resource management, agriculture, ecological modelling, ocean studies, etc. We propose to model this reasoning process as follows:

- **Interpretation phase:** Building a scenario of the situation through analysis of the spatial distribution of individual parameters given as sets of observation points and a model of the physical properties of each parameter.
- **Simulation phase:** Simulation of the evolution of the situation through application of physical models to the scenario. The simulation phase consists of a static and a dynamic part:
  - **Static inference:** Inference of non-observed parameters through combinations of observed parameter values.
  - **Dynamic inference:** Inference of the spatial evolution of parameters in terms of modifications to their spatial distributions.

The rest of this paper describes each of the above phases in turn, followed by a discussion where we put this work into perspective by comparing it to other approaches. We also discuss the utility of this approach and outline the current state of research and directions for future work.

2 Interpretation Phase

The goal of the interpretation phase is to build a scenario of the situation that can be used to detect which physical processes are causing the situation and simulate their evolution. In accordance with our study of meteorological practices, we propose to model the distributed parameters individually and use combinations of distributions to reason about the evolution of the system.

A physical system of spatially distributed parameters occupies a region of space, where each point can be assigned a value for each parameter. The values of each parameter are distributed in a specific pattern within the region. A *qualitative* description of this pattern is obtained by a double discretization: on the value domain of the parameter and on the space it describes. The value domain, e.g. the set of real numbers $\mathcal{R}$, is discretized into qualitative categories, e.g. intervals. An analogous spatial discretization is carried out on the points in the described space by grouping neighbouring points with equal values into larger spatial units, i.e. *regions*. The spatial distribution of the parameter is described qualitatively as a patchwork-like pattern of contiguous regions.

Figure 1 illustrates an example of the kind of physical system we want to model with this method.

The illustrated physical system is a cross-section of a part of the Earth–Atmosphere system, which can be described by different physical parameters, e.g. temperature, relative humidity, etc. In this scenario, the parameter *temperature* divides the space of the physical system into a pattern of three regions, corresponding to a discretization of the value domain $\mathcal{R}$ into the symbolic values {*cool warm hot* }. The parameter *relative-humidity*, on the other hand, divides the same space into a different pattern of only two regions, corresponding to its proper value domain discretization: {*dry humid* }.

The initial information on the distribution of a parameter is quantitative/metric and limited to the coordinates of the observation points and the values that have been observed at those points. Inferring the rest of the distribution from this sparse data requires a number of assumptions, which means that the resulting description will be qualitative in nature.

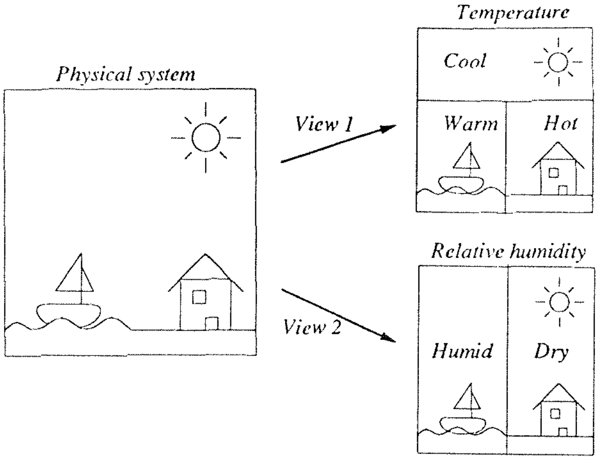

Figure 1: A physical system of spatially distributed parameters.

In order to distinguish between known quantitative/metric data and approximated/simplified qualitative data, we divide the description of the distribution of a parameter into two parts: a metric diagram and a place vocabulary. This division was proposed as a general model for qualitative spatial reasoning in [Forbus *et al.*, 1987]:

- The *metric diagram* describes the metric and quantitative properties of the world to be reasoned about. It is used for those queries that cannot be answered by purely qualitative reasoning.
- The *place vocabulary* describes the same world in qualitative terms.

Since the metric diagram and the place vocabulary describe the same world, although in different ways, they should be compatible. This is accomplished by using the quantitative information in the metric diagram to calculate the qualitative representation of the place vocabulary.

In this approach, the metric diagram consists of the set of observation points and a suitable connectivity structure. The place vocabulary is a topological map of the regions derived from the metric diagram. The construction of these two structures is supported by a model of the specific physical properties of the parameter. In the following sections, we will describe the construction and purpose of each of these components.

2.1 Metric Diagram: Connected Point Set

The goal of the analysis of the observation point set is to describe the spatial distribution of the parameter as a pattern of contiguous regions. This is accomplished by comparing the values observed at neighbouring points. If the same value is observed at two neighbouring points, then they can be considered qualitatively equal and grouped into a larger spatial unit, i.e. a region. If the observed values are not the same, then the two neighbouring points lie on the boundaries of two different regions.

Space is a continuous medium and consists of an infinite number of points. Between two points, there will thus always be an intermediate point, making the concept of *neighbour* very relative. Two points are only neighbours with respect to some level of approximation where all intermediate points are disregarded.

In the case of observation point sets, it is not always obvious which points are neighbours, since they can be spread out in an irregular pattern. Figure 2 shows a set of observation points for the parameter *temperature*. The observed values have been categorized into the qualitative values {*cool warm hot*}. The observation points can be in two or three dimensions, depending on which physical system is being modelled.

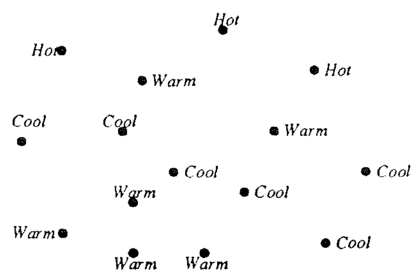

Figure 2: A set of observation points for the parameter *temperature*.

In order to know which points are neighbours and can be compared, a connectivity structure must be defined on the observation point set, i.e. a graph structure that indicates neighbourhood relations. A triangulation provides a natural connectivity structure for this kind of point set. A point set can be triangulated in many different ways. For the purpose of comparing observed values at neighbouring points, a triangulation that minimizes the distance between connected points is the most suitable. In [Preparata and Shamos, 1985], several algorithms are given for constructing various connectivity structures on point sets.

Figure 3 shows a triangulation of the point set in figure 2, where two points are neighbours only when the straight line connecting them does not intersect any shorter line connecting two points.

The metric diagram of our representation is the observation point set and a chosen connectivity structure. It is used to construct the place vocabulary, as described in the next section.

2.2 Place Vocabulary: Topological Map

The place vocabulary describes the qualitative, non-metric properties of the metric diagram. Whereas the metric diagram describes the spatial distribution of the parameter as a network of observation points, the place

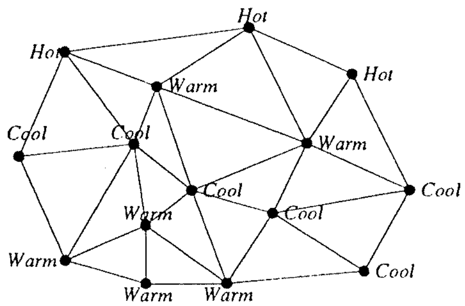

Figure 3: Metric diagram: triangulation of the point set in figure 2

vocabulary will describe the same distribution as a pattern of contiguous regions.

The place vocabulary is constructed by comparing neighbouring points in the metric diagram, according to the chosen connectivity structure, and grouping points with equal values into larger regions. Figure 4 shows how to detect regions in the metric diagram in figure 3.

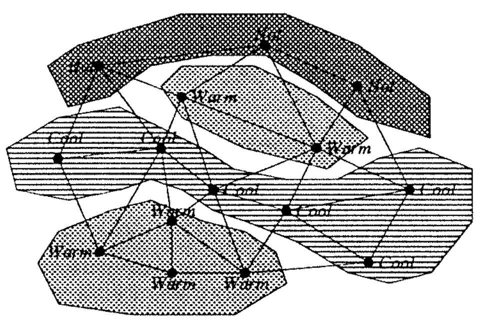

Figure 4: Regions detected in the metric diagram in figure 3.

The connectivity structure of the detected regions can be represented as a *topological map*. A topological map is an abstract illustration of the neighbourhood relations between regions, and does not convey any information on size or shape. Figure 5 shows the topological map of the regions in figure 4.

This particular topological map indicates the regions that can be detected by a straightforward analysis of the metric diagram. The next section describes how the topological map can be refined through the use of a model of the parameter's physical properties.

2.3 Parameter Model

During the interpretation phase, when representations of individual parameter distributions are being constructed, a model of the physical properties of a parameter enhances the information in the metric diagram and can lead to the inference of additional regions in the topological map or to a refinement of it.

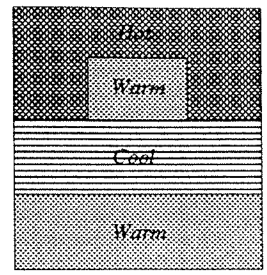

Figure 5: Place vocabulary: topological map of the regions in figure 4.

A parameter is defined by its name, unit and value domain, which can be finite or infinite. By dividing the value domain into different *quantity spaces*, i.e. sets of qualitative values, the distribution of the parameter can be described at varying levels of detail. Examples of quantity spaces are sets of intervals or symbolic values. The quantity spaces define alternative views of the value domain. The parameter model provides information on how to map between different quantity spaces and the value domain.

The value domain of a parameter is modelled as being either spatially *ordered* or *unordered*. This modelling choice depends on which properties of the physical system one wants to convey.

Spatially ordered value domains indicate that the spatial transition from one value to another must follow the order given in the quantity space and that there can be no discontinuities. This is a convenient property since it enables us to infer more information from the metric diagram than has actually been observed.

The topological map in figure 5 indicates that the two value regions *cool* and *hot* are neighbours. If the value domain of the parameter, in this case *temperature*, is defined as spatially ordered with the quantity space {*cool warm hot*}, then we can infer the existence of an intermediate *warm* region, although this value has not been observed. Figure 6 shows the resulting topological map.

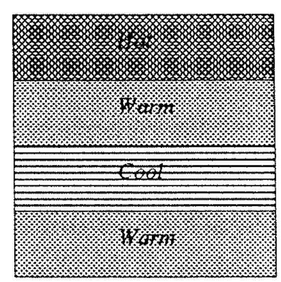

Figure 6: A refined version of the topological map in figure 5.

Parameters with spatially unordered value domains are equally common and have the property that any two values in the quantity space can correspond to neighbouring regions in the topological map. One example is the parameter *weather-type* which is an important observation in meteorology. An example of a quantity space for this parameter is {*rain thunder cloudy fog fine*}. Any two regions can be neighbours, thus it is not possible to infer any other regions than those detected in the metric diagram.

The metric diagrams and topological maps constructed for the distribution of each parameter make up a scenario describing the situation in a conceptual way. This scenario will be used during the simulation phase.

3 Simulation Phase

During the simulation phase, the constructed scenario is used to reason about the spatial properties and evolution of the physical system.

The scenario is a conceptual model of the situation, where the spatial distribution of each observed parameter is described by a metric diagram and a topological map. By modelling physical processes as topological maps and matching these against the topological patterns in the scenario, alternative descriptions of the physical system can be inferred and its evolution simulated.

The drawn inferences can be characterized as either *static* or *dynamic* as follows:

- **Static inference:** The scenario is used to infer the distributions of non-observed parameters through combinations of observed parameter values. This inference is static since it leads to alternative views of the physical system in the form of new parameter distributions, but existing parameter distributions are not modified.
- **Dynamic inference:** The scenario is used to infer the spatial evolution of existing parameter distributions, either observed or inferred. This inference is dynamic since it will modify the representation of existing parameter distributions. This may trigger further static or dynamic inferences.

The following sections will describe how to model static and dynamic inference respectively.

3.1 Static Inference

Reasoning about physical systems often means *combining* values of parameters in order to infer the value of some other parameter. A combination model describes which parameters are involved and how to calculate the result as a function of the parameters' values.

The combination model can be an equation or some other relevant computable function of several parameters. It can be expressed either as a qualitative version of an equation, e.g. using interval arithmetic, or as a matrix of value correspondences. Several versions of a combination model are possible to cater for all possible combinations of quantity spaces for the same value domain, i.e. different levels of granularity.

In the case of spatially distributed parameters, the values to combine must coincide in space as well as in time. As an example, consider the meteorological form of the equation of state:

$$P = \rho RT$$

P is pressure, ρ is density, R is the specific gas constant and T is temperature. P, T and ρ are distributed parameters with spatially ordered value domains, that can be discretized into intervals or symbolic values. R, the gas constant, also has a spatial distribution in the sense that it is applicable at all points where the specific gas has a distribution. P and T are readily observable parameters, whereas direct observation of ρ requires quite complicated equipment. It is thus convenient to infer the distribution of ρ from the given equation and the observed parameter distributions.

P and T are alternative views of the same region in space. By superimposing the spatial distributions of P and T, a new description of the same space emerges as a pattern of regions where the values of both P and T are constant. This pattern is the spatial distribution of ρ. The value of ρ in each region is calculated by applying the equation, or a qualitative version of it, to the values of P and T in these regions.

In the following sections, we will describe how to construct the spatial combination of two parameter distributions. We will also discuss how to handle the spatial ambiguity that may arise due to sparse data.

3.1.1 Combined Topological Maps

In order to infer the distribution of a parameter expressed as a function of other parameters, we must know which value regions intersect in space. For this purpose, a *combined topological map* is constructed for the involved parameters.

A topological map describes the connectivity structure, i.e. neighbourhood relations, of the regions *within* a single parameter distribution. Analogously, a combined topological map describes the topological relations *between* parameter distributions, i.e. where different regions intersect in space. The term for this topological relation is *overlap*. Two regions overlap if they have at least one point in common.

The metric diagram does not allow any inference of the exact shape of the different value regions. It is thus impossible to say exactly where and how two regions overlap. What can be inferred is whether two regions are *certain* to overlap, whether they *may* overlap or whether they are certain *not* to overlap. This amounts to finding out whether two regions have at least one point in common, do not have a point in common or it cannot be decided if they have a point in common.

The combined topological map is constructed by combining the metric diagrams of the parameters. In doing this, we want to infer which values could have been observed for the second parameter at the observation points of the first parameter, and vice versa.

In many practical applications, the parameters will have been observed at the same points, i.e. their metric diagrams will be spatially equal, only the observed values will be different. E.g. in the case of meteorology, most parameters are observed at the same observing stations. If two observation points are identical, then we know that the two regions, one for each parameter, that were inferred by means of that observation point are certain to overlap, since they have at least that point in common.

However, in the general case, two parameters need not have identical metric diagrams. By adding the observation points of the second parameter to the metric diagram of the first, we see that each new point falls within exactly one triangle in the metric diagram, as defined by the chosen connectivity structure. The values observed at the points connected by the triangle are the values that could have been observed at the newly inserted point.

Figure 7 shows an example of this situation. A point has been added to the metric diagram of the parameter *weather-type*, which has the spatially unordered quantity space {*rain thunder cloudy fog fine*}. The inserted point falls within a triangle connecting three observation points, where the values *rain*, *fine* and *thunder* have been observed. According to our assumption that value transitions take place between neighbouring points according to the chosen connectivity structure, exactly one of these values must have been observed at the inserted point.

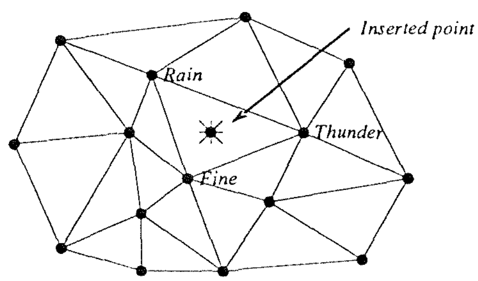

Figure 7: Inference of values that could have been observed at an inserted point.

This is an ambiguous situation with one, two or three alternatives, depending on how many different values have been observed at the three points connected by the triangle. The metric diagram does not allow us to decide which of the three regions the inserted point belongs to. However, it does allow us to decide whether a spatial intersection between regions in two different distributions is possible or not.

Figure 8 shows the different situations that can arise, assuming that the inserted point belongs to the metric diagram of the parameter *temperature* and the observed value is *cool*. The situations correspond to the following rules:

- **Certain overlap:** If two regions have at least one observation point in common then they are *certain* to overlap.
- **Possible overlap:** If two regions do not have any observation point in common, but some point falls within a triangle that has led to the inference of the region in the other distribution, then the two regions *may* overlap.
- **No overlap:** If the above rules do not apply, then the two regions are certain to be disjoint, i.e. they do *not* overlap.

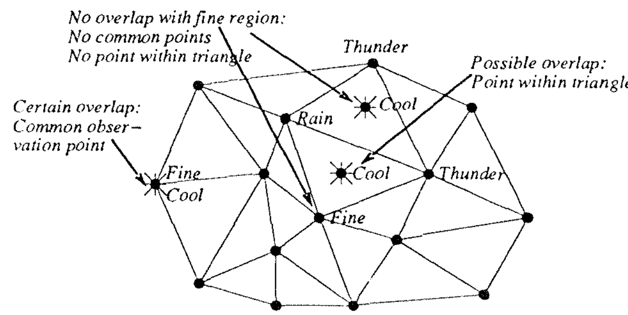

Figure 8: Inference of overlapping regions: the three situations deciding whether two regions are *certain* to overlap, *may* overlap or are certain *not* to overlap.

The combined topological map is constructed by comparing each pair of regions in the two topological maps according to the rules mentioned above. In case of ambiguity, the result is a *set* of combined topological maps, each indicating a possible overlap situation.

Figure 9 shows an example of two topological maps, for the parameters *temperature* and *weather-type*, and one possible combined topological map given the overlap structure indicated in table 1.

In the next section we discuss how to reduce the occurrence of spatially ambiguous situations.

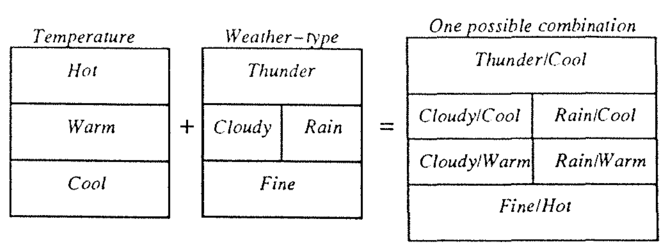

Figure 9: Topological maps for the parameters *temperature* and *weather-type* and one possible combined topological map given the overlap structure in table 1.

	Cool	Warm	Hot
Fine	Disjoint	Disjoint	Certain
Cloudy	Certain	Possible	Disjoint
Rain	Certain	Certain	Possible
Thunder	Certain	Certain	Possible

Table 1: Overlap structure for the topological maps in figure 9.

3.1.2 Controlling Spatial Ambiguity

In the previous section, we saw that combinations of topological maps sometimes contain ambiguous regions, where it cannot be decided whether two regions in the original distributions overlap or not. This ambiguity is due to the sparseness of data in the metric diagram. Human experts, e.g. meteorologists, use domain knowledge to disambiguate in this kind of situation. The following methods can be used to handle spatially ambiguous situations:

- **Treat the ambiguous region locally:** The ambiguity only concerns a pair of regions and is thus *local.* It does not influence the rest of the combined topological map, provided all other regions can be combined without ambiguity. Reasoning can thus continue unambiguously for a large part of the space of the physical system. The ambiguous region can be treated locally, either by indicating its value as unknown or by branching into multiple representations of that region.
- **Use proximity information to solve the ambiguity:** In some applications, proximity information can be used to disambiguate. A plausible model is to let an inserted point belong to the region of the observation point it is closest to. Figure 10 shows an example of this situation.

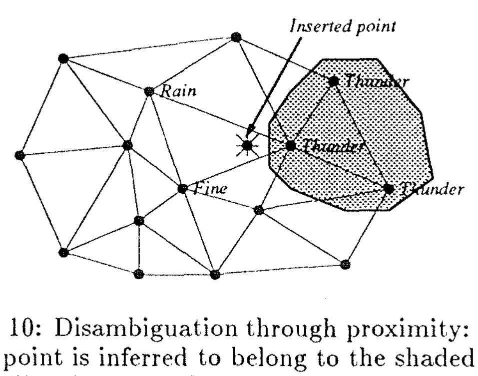

Figure 10: Disambiguation through proximity: the inserted point is inferred to belong to the shaded region since it lies closest to that observation point.

- **Use hierarchical parameter models to avoid unnecessary ambiguity:** Many parameters are physically relevant only in conjunction with some other parameter. By including this domain knowledge in the model, many potentially ambiguous situations can be avoided. Figure 11 shows an example of such a case. Again, the example is taken from climate modelling and illustrates a part of the Earth–Atmosphere system. The system is described by two parameters: *atmospheric-layer* and *soil-type.* The parameter *atmospheric-layer* divides space into regions according to the simplified quantity space {*stratosphere troposphere ground* }. The parameter *soil-type*, with the quantity space {*sand clay peat* }, can only describe points within regions described by the value *ground* for the parameter *atmospheric-layer.* Regions in the topological map of the parameter *soil-type* can thus only overlap with regions characterized as *ground* in the topological map of the parameter *atmospheric-layer*, and no further combinations need be considered in the construction of the combined topological map.

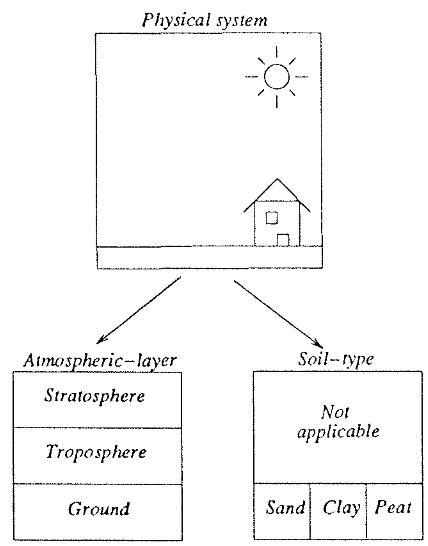

Figure 11: Hierarchical parameter models reduce the occurrence of spatial ambiguity.

3.2 Dynamic Inference

Dynamic inference differs from static inference in that it modifies the distributions of existing parameters instead of inferring new parameter distributions. In physics, dynamic evolution is usually modelled by differential equations. In qualitative physics, the temporal evolution of a parameter is usally modelled as transitions between subsequent landmark values in the quantity space of the parameter's value domain.

Analogously, in qualitative spatial simulation, the transitions will reflect significant changes to the spatial distributions of the parameters, or more precisely changes to their topological maps. Significant changes can take place either *within* a distribution, by rearranging the neighbourhood structure of the regions, or *between* distributions, in which case the overlap structure between regions is modified.

Physical processes are modelled as topological patterns of parameter regions that are matched against the scenario constructed during the interpretation phase. The spatial evolution of the system is given as a sequence of subsequent topological modifications to the parameter distributions.

Spatially distributed parameters often evolve through flow processes. We will outline a model of radiative flow from the sun through the layers of the atmosphere. Figure 12 illustrates the situation, which is, again, a part of the Earth–Atmosphere system, this time described by the parameters *emissivity*, *transmissivity* and *irradiation*.

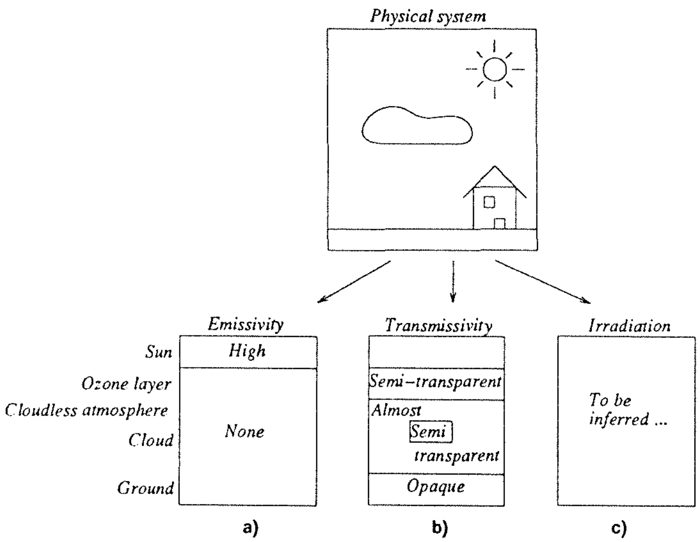

Figure 12: Topological maps for some key parameters in a model of radiative flow.

The parameter *emissivity* describes sources of shortwave radiation. Figure 12a shows the topological map of the parameter *emissivity* in this situation, using the quantity space {*high none* }. There is only one region of high emissivity, namely the sun.

The parameter *transmissivity* indicates how much of the radiation received by a region will be transmitted to more distant regions. The radiation that is not transmitted is either reflected or absorbed, which will increase the temperature of the region. However, those processes are not modelled in this example, which focuses on radiative flow. The topological map of the transmissivity regions in this situation is given in figure 12b. Its quantity space in this example is {*transparent almost-transparent semi-transparent almost-opaque opaque* }. There are four different transmissivity regions in the scenario, corresponding to the *semi-transparent* ozone layer, which filters a lot of the radiation coming into the atmosphere, the *almost-transparent* cloudless atmosphere, a *semi-transparent* cloud, and finally the *opaque* ground, which absorbs or reflects all radiation it receives and transmits none. The transmissivity of the sun is not relevant to this model, so we leave the corresponding region unspecified.

The parameter *irradiation* indicates regions that receive radiation. The initial distribution of this parameter in figure 12c indicates no irradiated regions. The model will describe the spatial evolution of the distribution of this parameter. The final distribution will indicate which regions in space receive more radiation than others.

In this model, we want to reason about how some regions are shadowed by others, and thus receive less radiation. In order to do this, it is necessary to include the notion of *flow direction* in the model. *Direction* is a spatially distributed parameter that divides the space of the physical system into qualitative vector fields with respect to some region. Figure 13a shows the distribution of the parameter *direction* with respect to the sun, i.e. the *high* emissivity region in figure 12a. Figure 13b shows another distribution of *direction*, this time with respect to the ozone layer, i.e. the upper *semi-transparent* transmissivity region in figure 12b. Finally, figure 13c shows the distribution of *direction* with respect to the cloud, i.e. the smaller *semi-transparent* transmissivity region in figure 12b. The value domain has been discretized into the categories {*inside beneath above left right* } which are convenient to this model.

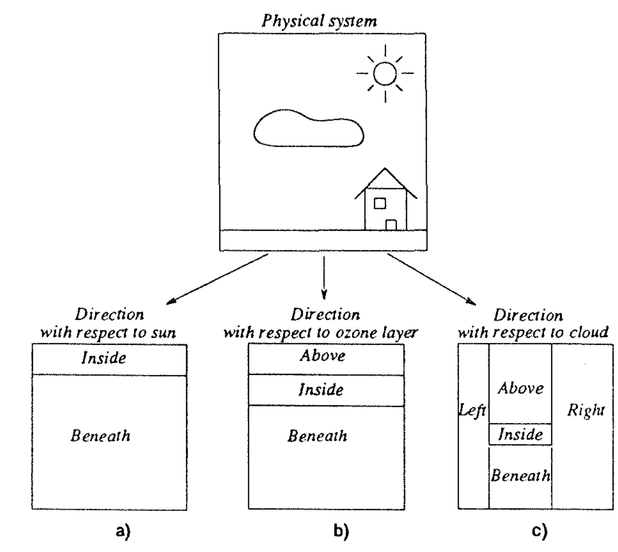

Figure 13: Topological map of the metric parameter *direction* with respect to different regions.

The simulation of radiative flow proceeds in the following steps:

- Regions of high emissivity match the pattern required for the physical process *radiative-flow*. There is only one region that matches this description in the topological map of the parameter *emissivity* (figure 12a), namely the sun. The region of ***high*** emissivity becomes the source region of this instance of the flow process.

- Once a source of radiative flow has been found, the regions that it will flow into must be detected. Flow follows a spatial order, so the receiving regions will be neighbours of the source region. The source region has only one neighbour in its topological map (figure 12a), which is indicated by the value *none*. This becomes the sink region for the flow.
- Once there is a source region and a sink region, the direction of the flow can be determined. Figure 13a shows the distribution of the parameter *direction* with respect to the source region, i.e. the sun. The overlap structure between this topological map and that of the parameter *emissivity* indicates that the sink region is totally contained within the *beneath* region. This value becomes the direction of the flow.
- The overlap structure between the topological maps *emissivity* and *transmissivity* indicates that the space occupied by the designated sink region contains several different regions of transmissivity. The flow will proceed gradually through these regions.
- The first step is the region of *semi-transparent* transmissivity that lies closest to the source, i.e. the ozone layer. That region will receive all the radiation sent from the emitting region, i.e. 100%. This is indicated in the model as a modification to the topological map of the parameter *irradiation*. A region that corresponds to the ozone layer is introduced into the *irradiation* distribution and given the value *100%*, see figure 14a. For the sake of this example, we will not bother with defining a quantity space for the parameter *irradiation*, but simply indicate the irradiation with approximative percentages.
- The flow will pass through the irradiated region according to the inferred flow direction. However, only a part of the received radiation is transmitted, as some of it is absorbed or reflected. Consultation of the overlap structure between the parameters *irradiation* and *transmissivity* indicates how much radiation will be transmitted. In this case, the irradiated region corresponds to a region of *semi-transparent* transmissivity, so we presume that 50% of the radiation passes through it. In the real model, a suitable qualitative equation would be used.
- The flow from the current irradiated region, i.e. the ozone layer, proceeds into neighbouring regions of constant transmissivity. The flow parameter indicates that these regions must also lie *beneath* the initially irradiated region. This is the case for the cloudless atmosphere, indicated by *almost-transparent* in figure 12b.
- However, the flow model also requires that the receiving region have no holes. The topological map in figure 12b indicates that the *almost-transparent* cloudless atmosphere contains a region of lower transmissivity, namely a cloud. The correct region to irradiate is constructed by removing the cloud and the area *beneath* it, according to the flow parameter, from the cloudless atmosphere. The resulting irradiated region is shown in figure 14b. Its value is indicated as *50%*, reflecting that some of the radiation was absorbed by the preceding region in the flow.
- The radiation continues to flow through the atmosphere, reaching the cloud, the shadowed region beneath the cloud and the ground. The final distribution of the parameter *irradiation* is shown in figure 14c.

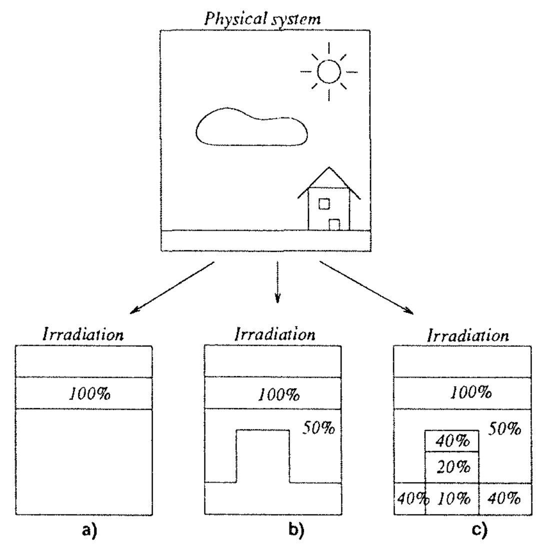

Figure 14: Steps in the inference of the irradiation distribution.

In this model, the existence of a cloud resulted in non-uniform radiation of the ground. The received radiation will be absorbed and transformed to heat according to the distribution of e.g. the parameter *heat-capacity*, thus creating a pattern for the parameter *temperature*. Differences in temperature often trigger other physical processes, e.g. sea breezes, cloud formation, plant growth, etc. These are just a few examples of physical processes that can be modelled with spatially distributed parameters.

4 Discussion

This paper has described work on modelling the qualitative behaviour of physical systems of spatially distributed parameters. We will put this work into perspective by comparing it to some other approaches.

In qualitative physics research, physical systems are often modelled as sets of spatially discrete objects.

The parameters are seen as attributes that describe the objects they are associated with. The model describes how the objects interact through their parameters. Object-oriented models are appropriate for many applications. see e.g. the thermodynamics model in [Collins and Forbus, 1991].

However, in applications like meteorology, it is not evident which objects the model should be built on. The system is more appropriately described by the spatial distribution of the individual parameters. The traditional approach in physics is to model a physical system as a set of differential equations, each describing how the value of a parameter varies as a function of some spatial axis. By combining different dimensions as needed, a full description of the parameter is obtained. Work on using differential equations in spatial qualitative resoning has been presented e.g. in [Throop, 1989] and [Nishida, 1993].

A third approach, based on topology, is given in [Cui *et al.*, 1992] where the process of phagocytosis, i.e. amoeba ingesting food, is modelled as a sequence of topological relations between discrete regions in space.

Modelling a physical system in terms of objects is attractive, since it gives the model a tangible touch. It is appropriate both for device- and process-oriented qualitative simulation and the envisioned situations can easily be illustrated diagrammatically. However, these models fail to convey the notion of continuous spatial distribution, which is readily modelled by differential equations. Models based on explicit differential equations are, however, not as readily understood by non-experts and are also not available for all kinds of physical systems.

The work presented in this paper can be seen as a combination of these two modelling approaches, where the spatial distributions of parameters are divided into patterns of discrete regions that can be manipulated as objects.

The utility of qualitative models of spatially distributed physical processes is manifold. Such models would provide a reasoning component for Geographic Information Systems (GIS) and programs for scientific visualization. They would provide a means of communication between professionals, e.g. meteorologists, by making it easier to hand over analyses of spatial situations to the next person on the shift. Their utility for pedagogical purposes is obvious. In fact, most of the examples in this paper have been taken from textbooks on meteorology and climate modelling, which, although their main purpose is usually to convey a quantitative understanding of the atmosphere, often devote a substantial part of each chapter to qualitative and diagrammatic descriptions of atmospheric processes.

Work on this approach continues actively. The next step will be to refine and implement the methods described in this paper. Models of basic atmospheric processes, such as radiation, conduction, convection and advection, are being developed, and will be integrated in a model of a fairly complex atmospheric process: the life-cycle of a sea breeze. We are also investigating applications in agriculture and natural resource management.

5 Acknowledgments

I would like to thank the Swiss National Science Foundation for supporting this research under SPP Project No 5003-034269. Mr. Jacques Ambühl at the Swiss Meteorological Institute has kindly introduced me to practical weather prediction and pointed me to relevant literature. I would also like to thank Professor Boi Faltings and Dr. Dean Allemang for useful comments.

References

Collins, John and Forbus, Ken 1991. Building qualitative models of thermodynamic processes. Technical report, Beckman Institute, University of Illinois.

Cui, Z.; Cohn, A. G.; and Randell, D. A. 1992. Qualitative simulation based on a logical formalism of space and time. In *AAAI*. 679–684.

Forbus, Ken; Nielsen, Paul; and Faltings, Boi 1987. Qualitative kinematics: A framework. In *IJCAI*. 430–436.

Nishida, Toyoaki 1993. Generating quasi-symbolic representation of three-dimensional flow. In *Seventh International Workshop on Qualitative Reasoning about Physical Systems.*

Perby, Maja-Lisa 1988. Computerization and skill in local weather forecasting. In *Knowledge, Skill and Artificial Intelligence*. Springer-Verlag. 39–52.

Preparata, Franco and Shamos, Michael 1985. *Computational Geometry: An Introduction.* Springer-Verlag.

Throop, David R. 1989. Spatial unification: Qualitative spatial reasoning about steady state mechanisms, an overview of current work. In *Third International Workshop on Qualitative Reasoning about Physical Systems.*