

Diagrammatic Reasoning about Actions

Using Artificial Potential Fields

Marcello Frixione[‡], Gianni Vercelli^{*}, Renato Zaccaria[†]

[‡] Dept. of Communication Sciences
University of Salerno
Via Ponte Don Melillo
I-84084 Fisciano (Salerno), Italy
frix@dist.unige.it

^{*} DEEI - University of Trieste
Via A. Valerio, 10
I-34127 Trieste, Italy
vercelli@univ.trieste.it

[†] DIST - University of Genova
Via Opera Pia, 13
I-16145 Genova, Italy
renato@dist.unige.it

Abstract

In this paper we propose a diagrammatic framework, as a way of performing spatial reasoning about dynamic scenarios. Our proposal is based on intrinsically dynamic models, that act as some sort of mental simulation, a "mental movie" of the evolution of the represented world. This simulation is generated starting from some reference "pictures", or "snapshots" of the represented situation (that we call *icons*). Icons are interpolated by suitable generative processes, and they drive the dynamic simulation. The mechanism that interpolates icons in mental simulations is based on the model of *artificial potential fields* (APFs). APFs are a well-known method used for motion planning in Robotics. We show that the use of APFs can be extended to different domains, in order to face various problems of spatial reasoning and planning. In the paper we present our method by describing and discussing a series of examples, concerning both spatial reasoning and the interaction of reasoning with perception.

Introduction

The role of image-like representations in the computational modelling of mental processes is a long-standing topic in the fields of Artificial Intelligence and Cognitive Sciences. Many authors proposed some kind of pictorial representation (in a broad sense) in order to account for human cognitive abilities and/or to design "intelligent" computer programs. In the early years of AI, diagrams played an important role in Gelerntner's theorem proving machine (Gelerntner 1959). Another influential example in AI has been the use of diagrammatic representations in problem solving proposed by Funt (1980). In the field of cognitive psychology, Kosslyn and other authors developed accounts of mental imagery in terms of pictorial representations (Kosslyn 1980).

In the past, image-like representations have often been considered sharply separated from more linguistically oriented approaches, and, in some sense, opposed to them

In the context of the debates between supporters and adversaries of the use of logic in AI, image-like representations were generally contrasted to logic based formalisms. Within the field of the Cognitive Sciences, the opposition between propositional and pictorial explanations of mental images involved psychologists and philosophers in lasting debates (see e.g. Block 1981).

In more recent years we assisted to a renewed interest in pictorial and diagrammatic representations in AI (witnessed for example by Chandrasekaran and Simon 1992; Kulpa 1994; Glasgow *et al.* 1995). If compared with the debates of the seventies, this revival is characterised by a more ecumenical mood. Researchers do not aim to individuate a clear-cut distinction between two classes of representations (propositional vs. pictorial or similar), nor they believe that such a crisp boundary exists. A rich spectrum of different types of representations has been identified, that share some kind of feature with images and diagrams, and that are more or less closely related to propositional systems. At the one extreme, there are logic-based representations that can be considered "analogical" in the sense of being isomorphic to their logical models (as is the case of vivid KBs in the sense of Levesque 1988). A similar approach is shared by symbolic systems that are not logic-oriented, such as mental models (Johnson-Laird 1983) or mental spaces (Fauconnier 1985). In other proposals, image-based representations are conceived as closer to the output of some level of the perceptual systems. Hybrid models have been developed, in which logic-oriented or rule-based modules interact with diagrammatic representations (Forbus 1995; Myers and Konolige 1995). In many cases the differences between image-like and propositional representations are seen in terms of computational advantages, rather than a matter of expressive power (Sloman 1995).

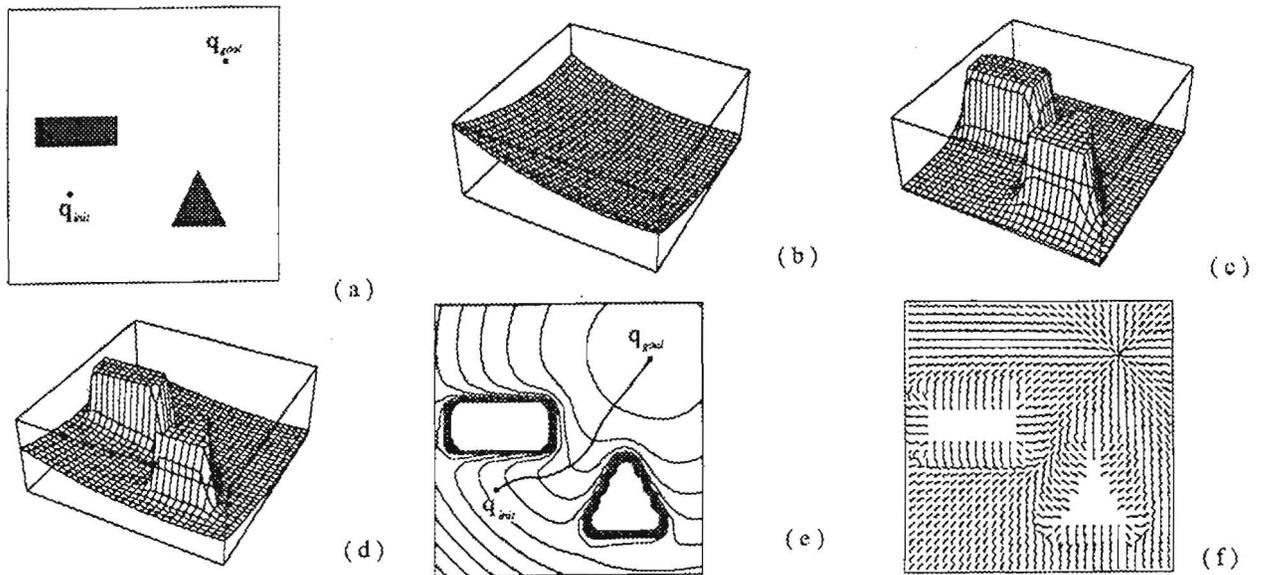


Figure 1 – A simple APF

In this paper we present a framework based on diagrammatic representations for reasoning about *dynamic* aspects of the world. Our proposal is based on intrinsically dynamic models, that act as some sort of mental simulation, a "mental movie" of the evolution of the represented world. This simulation is generated starting from some reference "pictures", or "snapshots" of the represented situation (that we call *icons*). Icons are interpolated by suitable generative processes, and they drive the dynamic simulation. In this sense, icons act as the key postures of the cartoon terminology. The mechanism that interpolates icons in mental simulations is based on the model of *artificial potential fields* (APFs), which are described in the following section.

Our approach has been developed according to the ecumenical spirit described before. We do not maintain that some qualitative and irreducible difference exists between image-like and propositional representations. We simply assume that our *representations* are akin to *pictures* in the sense that:

- They are spatial representations where information concerning space is represented in an implicit rather than in an explicit way. Similarly, our representations are intrinsically dynamic, and temporal aspects are represented implicitly by the temporal evolution of the model.
- They are made out of primitive elements that do not correspond to high level entities: the atomic constituents of the representations are not, in their turn, representations at the level of macroscopic objects of the common sense world. Rather, the system is based on more fine-grained spatial

primitives. In the examples of the present paper we assume that the primitive elements of our iconic models are similar to the pixels of a bitmap.

- They represent the evolution of specific states of affairs, and do not allow the representation of pieces of incomplete knowledge (e.g. disjunctive knowledge, or quantified knowledge).

These assumptions offer various advantages if contrasted with more traditional approaches to action representation. For example, in the domain of spatial reasoning, our model exploits the rich source of geometric information implicitly encoded in pictorial representations, thus avoiding the problem of working out exhaustive explicit descriptions of shapes and their spatial arrangement. Moreover, the pictorial format of spatial representations makes them closer to the data coming from sensors. This makes easier the interaction between perceptual and reasoning tasks in autonomous systems. The choice of bitmap-like representations is not crucial. A more structured spatial representation system based on some kind of geometric primitives would work equally well. What is essential is that the representations convey richer geometrical information if contrasted with traditional propositional systems

In the following, we introduce the model of APFs, and show in which sense they can be used for the interpolation of icons. We also present the main drawback of APFs, i.e. *local minima*. Then, we discuss the use of APFs for reasoning on dynamic diagrams and for planning complex tasks. Successively, we propose APFs as a way to link together perceptual data and diagrams. Finally, we show how perception and diagrammatic reasoning can be merged in an example of on-line reconstruction of a

complex dynamic scene.

Artificial Potential Fields for Diagrammatic Reasoning

The *Force Field Metaphor* became popular in Robotics, Automatic Control, and Artificial Intelligence after a paper by Khatib (1986), in which this model was originally used for generating trajectories of an articulated robot in presence of obstacles. The basic idea is simple: to build a global *force field function* over the world at hand, including moving entities and obstacles, in which obstacles are sources of *repulsive forces* and the goal generates an *attractive force*. Once the field is set, and a dynamic model for the moving entity has been chosen, the motion is solved simply by “observing” the motion of the entity subject to the effect of the field forces. It is worth noting that the field itself (as well as the dynamics of the moving entity) needs not to refer to the physical properties of the problem, hence it is purely metaphorical, or abstract. This is true also for the kind of *field function*, which is only required to have a (usually scalar) *potential function*; it is often inspired to gravitational or electric fields. For this reason the method is usually known as *Artificial* (or *Abstract*) *Potential Field* (*APF*). The example of fig. 1 (a) to (f) is taken from (Latombe 1991) and shows the simplest use of an APF for 2-D navigation of an ideal robot towards a target \mathbf{q}_{goal} starting from \mathbf{q}_{init} in presence of two obstacles (fig. 1a). A global artificial potential function is built as the sum of an attractive one (“pulling” toward \mathbf{q}_{goal}) (fig. 1b) and a repulsive one (“pushing” away from obstacles) (fig. 1c). The resulting APF is shown in figures 1d and 1e; in the latter the resulting trajectory is also shown, generated by a “free motion” of the robot descending the field’s gradient. Figure 1f shows the force field associated to the APF. Note that $U(\mathbf{q}_{\text{goal}})$ is a *global minimum* of the field. For each obstacle B_i , a typical definition of the repulsive part of the APF at any position \mathbf{p} is:

$$U_{B_i}(\mathbf{p}) = \begin{cases} \eta \left(\frac{1}{\rho_i(\mathbf{p})} - \frac{1}{\rho_0} \right)^2 & \rho_i(\mathbf{p}) \leq \rho_0 \\ 0 & \rho_i(\mathbf{p}) > \rho_0 \end{cases}$$

where $\rho_i(\mathbf{p})$ is the minimal distance between the robot and the obstacle, and ρ_0 is the spatial limit of influence of the obstacle. Similarly, the attractive component, due to the goal, is $U_{\text{an}}(\mathbf{p}) = \frac{1}{2} \xi \rho_{\text{goal}}^2(\mathbf{p})$ where $\rho_{\text{goal}}(\mathbf{p})$ is the distance from \mathbf{q}_{goal} ; η and ξ are suitable constants. Hence, the overall APF at any point \mathbf{p} is defined by the function $U(\mathbf{p}) = U_{\text{an}}(\mathbf{p}) + \sum_i U_{B_i}(\mathbf{p})$. This field, after suitable differentiation, defines the artificial force acting, at any instant and at any \mathbf{p} , on the robot; given a simple dynamics, this generates a plausible motion law. In this example, a

conservative potential function $U(\mathbf{p})$ is used; in this way the equipotential lines are closed and unique. This guarantees that, for each position \mathbf{p} , there exists a unique possible motion towards \mathbf{q}_{goal} . Note how it is sufficient an almost local knowledge (the minimal distance $\rho_i(\mathbf{p})$ for each obstacle B_i) to set up the field and to drive the motion at any instant.

The example in fig. 1 has only two dimensions. Commonly, robots with more degrees of freedom (e. g., a robotic arm with a hand) need a multi-dimensional representation in which the APF approach is successfully used as well.

The APF approach introduced several novelties. First, it faces successfully the problem of generating real-time motion by mixing *planning* and *control* concepts. Second, it is based mostly on *local knowledge* about the world, by measuring few geometric quantities (i.e., distances) around the moving entity, using them to set up a global field function, and generating the motion as a trajectory descending the field; the use of local knowledge makes the method computationally tractable. Third, it makes no difference, in principle, between static and dynamic (evolving) worlds, since every set of measurements can at any time setup (or modify) the global field. Lastly, the force field can be built over a multi-dimensional space, in which each dimension is a different degree-of-freedom of the (moving) system itself. For example, in the case of a common 6-axes robotic manipulator, the space (usually known as *configuration space* or *C-space* (Lozano-Pérez and Wesley 1979)) has dimension 6. Although not very common in robotic literature, a multi-dimensional space might also model and plan multiple moving entities in the same world, avoiding each other (Latombe 1991).

Historically, APFs were studied in earlier works about the control system of living beings (humans and mammals in particular); the force-field concept has gained central stage after the seminal paper by A.G. Feldman (1966) and, among controversies, is still a significant factor which has even widened its breadth. Originally, the concept was limited to the physical domain of muscle elasticity, which implies a potential energy function and hence equilibrium points or stable postures. The idea was then extended to trajectory formation, interpreted as a dynamic process, which brings the initial posture to a final planned one by following the flow-lines of the corresponding force fields (Morasso *et al.* 1997). It was in successive works (Loeff and Soni 1975, Khatib 1986, Connolly *et al.* 1990) that the force fields became a computational metaphor popular in Robotics for expressing constraints and incorporate them in trajectory formation algorithms, with the interesting characteristic of working mostly on *local* instead of *global knowledge*.

Local Minima

It was soon recognised that mixing attracting and repulsing

components would result in global fields with local minima, thus involving deadlocks. In effects, a simple navigation problem can be solved by letting the robot descend the gradient of the field. However, in many cases the direction of the gradient does not lead to the goal (the global minimum of U), and the robot gets trapped in a *local minimum* of U . Consider, for example, a navigation situation (figure 2) in which A is the starting position, and B is the goal position. A V-shaped obstacle gives rise to a local minimum Lm. The gradient descent would lead the robot in Lm. Several methods have been proposed to overcome this kind of problems, most of them with a markedly heuristic character; nevertheless, these methods cannot be easily applied unless a global knowledge of the environment is available (Koditschek 1987; Connolly *et al.* 1990; Koren and Borenstein 1991; Tilove 1990).

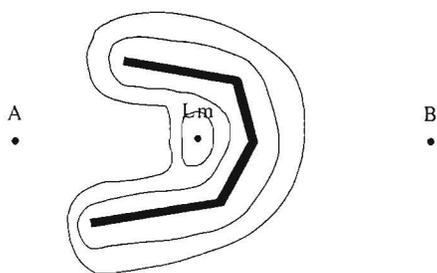


Figure 2 – A local minimum

In general, the power of the various kinds of force-field methods lies in the fact that, by definition, a single flow-line passes through any point in the working space, thus guaranteeing unique solutions that do not need decisions and/or reasoning. This, however, limits the richness of possible behaviours and, in some cases, due to local minima, it is not sufficient to reach the goal. Therefore, in general, the need emerges for a computational framework, in which APF-based methods are combined with some form of discrete global process, often of a heuristic nature. APFs are still the computational backbone of the system, and generate the continuous outflow of motor commands. In addition, such foreground process is modulated by background discrete processes that carry out a variety of simple cognitive functions (classification of situations, detection of periodic events, local distortion of the fields). This has lead, among others, to models based on *schemata* (Arkin 1990) or navigation *templates* (Slack 1993). Informally, a schema could be defined as an heuristic that allows to modify the “natural” evolution of an APF, in order to avoid local minima.

In Robotics APFs are widely applied to solve the *path-planning* problem in robot navigation. Local minima require suitable backtracking techniques. In multi-dimensional spaces (e.g., while modelling a manipulator in the *configuration space*) these techniques may become complex and computationally heavy, but are nevertheless a well-settled and currently used methodology (Latombe

1991).

Mental Simulation of Dynamic Worlds based on Diagrams

The example of fig. 1 could be considered as a simple form of diagrammatic reasoning. The APF mechanism operates on a picture-like representation of the scenario. In addition, even the inference mechanism is, in some sense, “diagrammatic”, in that the solution of the problem is achieved through an evolution of the model that mimics a possible evolution of the represented situation. Our proposal is to extend this type of approach to other forms of reasoning about actions and dynamic scenarios. Roughly speaking, the general idea is that goals can be modelled as attractive targets and constraints (of different kinds) can be considered as repulsive obstacles in a multi-dimensional space. The initial configuration in a n-dimensional space (as, for example, figure 1a) can be seen as a diagram, in which all information is present to set up a field. Then, we can look at “how the field evolves” waiting for some solution of the problem. In this way, APFs can be adopted as a way of representing/supporting *diagrammatic simulations* for spatial reasoning in dynamic scenarios. In particular, this way of reasoning can be used to plan actions in spatial contexts, and in presence of constraints that can be in general represented as repulsive “obstacles”. The techniques developed in Robotics in order to avoid local minima are usually oriented to on-line planning and control rather than to reasoning tasks. On the contrary, in the following we shall focus on methods that are suitable for reasoning. In this context, APFs’ main drawback (local minima) becomes less important, while some other characteristics that are not relevant in normal planning/control problems in Robotics become very appealing. Technically, this is true every time APFs are used to solve “quasi local problems”, as to say, in which *the goal is rather close to the actual position* in the state space. This statement seems obvious and discouraging: how useful is a planner capable of solving problems in which the goal is near the start? However, we first must note that this vicinity is in terms of “difficulties” encountered, not in terms of geometric distance (e.g., Euclidean metrics). Furthermore, if we have memory enough, we can use a general task representation in terms of partially ordered diagrams, each of them representing a “goal” or a “constraint”, while an APF supports “local planning” in each sub-plan. This form of representing plans is appealing, because it can use expressive diagrammatic representations, and allows plan classification, composition and adaptation.

The spirit of the approach can be understood considering a fragment of an instruction sheet (Bosch Hausgeräte GmbH 1990) for installing a dishwasher (fig. 3). Note how the assembly task is represented as a partial ordering of

diagrams; each diagram is either a “sub-goal” or a “constraint” (the latter are marked with a cross, to mean “impossible” or “stay away from it”). Performing the task is a form of plan adaptation involving an amount of (local) planning. This kind of representation is plausible, high-level, and, at least in principle, it is consistent with an APF approach.

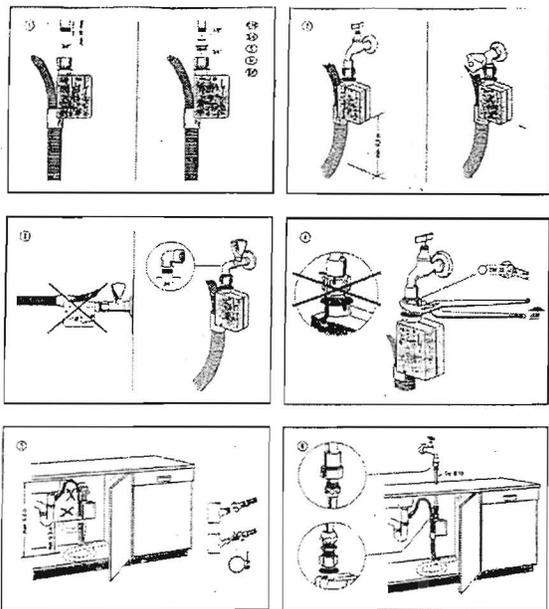


Figure 3 - Instruction sheet for installing a washing machine

Figure 3 shows that, in order to accomplish the task, one must reach the goals depicted in frames 1 to 6 (in this case, they are totally ordered), “staying away” from states depicted in frames 3-left and 4-left, which have a partial ordering relation with other frames.

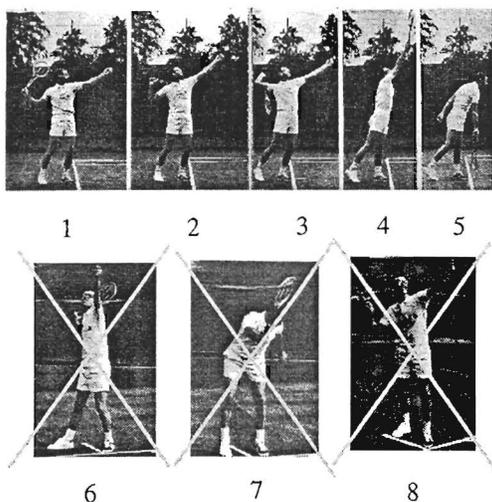
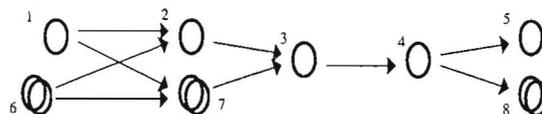


Figure 4 – A tennis example

Another evocative example is shown in figure 4. It is a picture taken from a tennis handbook (Clerici 1972) that shows how to serve. Here a pure kinematic task is described in a way similar to the example of fig. 3. A sequence of correct postures (frames 1 to 5) is interleaved with wrong postures to be avoided, that are shown at the bottom of the figure (frames 6-8).

We generalise the use of partially ordered diagrammatic descriptions like the ones above, and think of multi-dimensional “snapshots” of reality (including kinematic descriptions of the world as well as other state variables like force, temperature and so on). In this paper, we call *icons* such representations. APFs act as the “inference engine” interpolating the icons in the diagrammatic simulation.

It is worth noting that in this form representation time is intrinsically implicit. Time appears when, in the mental simulation, we go from one icon to another while executing the task. Using a straightforward graphic representation, we can depict the task of figure 4 as follows:



This graph corresponds to a (partial) plan. Ellipses indicate icons, double ellipses “negative icons”, arrows ordering relations. The semantics of negative icons is not only related to generic “situation to avoid”, but also to “deadlocks” (usually, local minima) in the related APF-based planning. In synthesis, *positive icons* represent goals and sub-goals, and act as *attractors* in the APF; *negative icons* represent *obstacles*, and, more in general, constraints of different kinds, and act as *repulsors* in the APF.

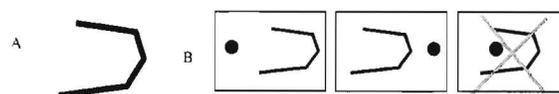
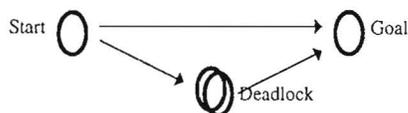


Figure 5 - Icons describing a navigation problem

This kind of icon-based representation can be thought as a generalisation of acting/planning models. At the one extreme, a planning problem can consist only in the initial and final icons; at the opposite, an action can be described in details by means of a rich set of (partially) ordered icons. Consider for example a mobile robot travelling from A to B in presence of a concave obstacle (figure 5, left). A description of the task may be based only on the Start and Goal icons, asking the APF to solve deadlock problems. Alternatively, it may be based on the three icons of fig. 5 (right), where the third icon helps the system in solving the problem. Note that the third icon is a negative one. The

ordering relationship between the three icons can be the following:



This can be considered an example of a *schema* in the sense of (Arkin 1990) as mentioned before.

In many cases negative icons can be “discovered” during the mental simulation itself in a surprisingly simple way. Consider the example of figure 6. On the left it is shown the initial icon of a navigation across a V-shaped obstacle, from P (bottom) to G (top). The field lines of the APF are also shown.

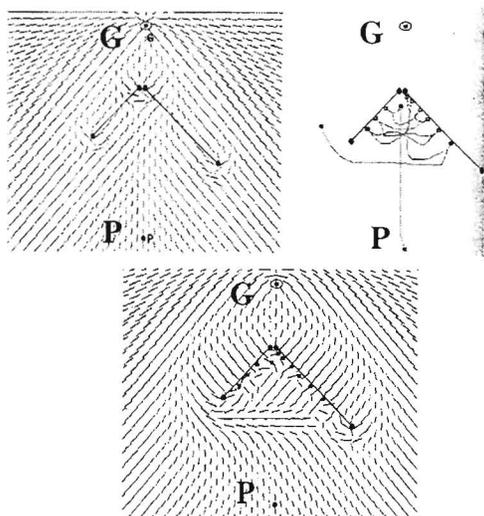
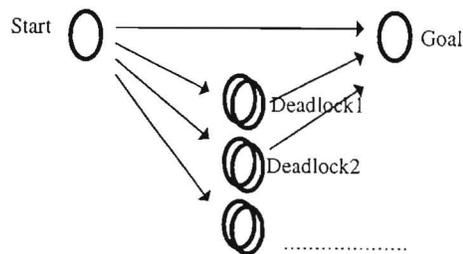


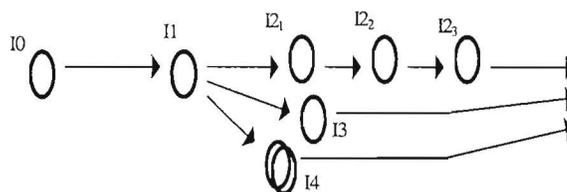
Figure 6 - Automatic discovery of negative icons

Due to the concavity, the robot would get trapped inside the obstacle (where the potential field U exhibits a local minimum). A way to overcome the problem would be the introduction of a negative icon similar to that of figure 5. However, the problem arises of explaining how (and by whom) such an icon could be generated. A plausible solution consists in the strategy of iteratively put a repulsive point into every discovered local minimum: every time it reaches a deadlock, it generates a “negative icon” acting as a new constraint (fig. 6, upper right). On figure 6, lower, are shown the field lines resulting from the superimposition of all the negative icons autonomously “discovered” during navigation. Hence, from this point of view, *negative icons* can be considered as a sort of “virtual obstacles” that are added to the mental simulation. They act as “temporary data” in iconic reasoning. The resulting ordering graph for the problem of figure 6 is a generalisation of the previous one:



This interactive process is a form of diagrammatic dynamic reasoning.

An appealing feature of such dynamic, icon-based framework is the possibility of merging different representations in a natural way by simply adding new icons. This compositional character is difficult to achieve using traditional deliberative approaches, and it is made possible by the fact that APFs are intrinsically compositional: adding new icons to a representation the evolution of the dynamic model changes in an incremental way. This can be used also for adding constraints. Consider the task “turning a screw using a screw driver”. You must align the blade with the slot on the screw head (I0), then put the blade inside the slot (I1), then turn (I2) while performing a given pushing force (I3) and avoiding misalignments (I4). The graph of a natural diagrammatic representation based on icons could be the following:



Note that “turning” is represented as a sequence of icons, and I3 has “Newtonian” dimensions (including dynamics) other than usual Cartesian, kinematic co-ordinates. I4 is also a representation of a multi-dimensional obstacle that prevents the screw driver to get mis-aligned (a sort of funnel surrounding the hand-held tool). The examples discussed so far show the different roles played by negative icons or, more in general, by repulsing entities in the APF: obstacles, generically unwanted situations, virtual obstacles (to prevent deadlocks), constraints.

APFs for Reasoning about Icons

When used to solve the “quasi local problem” of going from one icon to another, APFs constitute a form of *analogical reasoning* about dynamic, spatio-temporal problems. As mentioned above, local minima can be avoided thanks to the short distance between icons, and thanks to the presence of negative icons repulsing the model from local minima. Furthermore, negative icons can be generated by the mental simulation itself during analogical reasoning. So, the principal drawback of APFs

(local minima) does not constitute a problem. On the contrary, local minima are a source of information comparable to *backtracking points* during discrete search within a search space. Another favourable point is the intrinsic capability of composing the effects of icons, by simply superimposing the induced fields.

The APF acts hence as an “interpolating engine” which generates an action that satisfies a sequence of multiple icons, keeping away at the same time from *physical* obstacles, as well as from *virtual* obstacles corresponding to unwanted situations, or constraints. It is worth noting that this generative process is based on an energy metaphor: an energy function is defined for every point of the field, and the process of finding a path from a start icon to a goal icon is a process of *relaxation* from a higher energy level to a lower energy level. Everything happens in the APF might as well be described in terms of energy values: the start icon is a local maximum of energy, the goal icon is a global minimum. Negative icons are obstacles, that can be used to prevent the model to get trapped inside a local minimum.

A more complex example is discussed in (Ardizzone *et al.* 1993) and can be summarised here. Consider an initial situation (represented by the icon in fig. 7a), in which a cube, a box and its cover are randomly placed on the ground. In the goal situation, represented by the icon in fig. 7b, the cube has been put inside the box, and the box has been closed by the cover. In order to plan a correct sequence of actions, an APF is set, in which the cube and the cover are both attracted towards their final positions. In this initial version a deadlock situation can occur, in which the cover reaches its goal position before the cube has entered the box (fig. 7c). Situation 6c can be recognised, while reasoning in the mental model, to be a local minimum in the field. The deadlock is then avoided by adding 6c as a negative icon to the APF, which constrains the evolution of the model, in order to obtain the correct sequence of actions (fig. 7d).

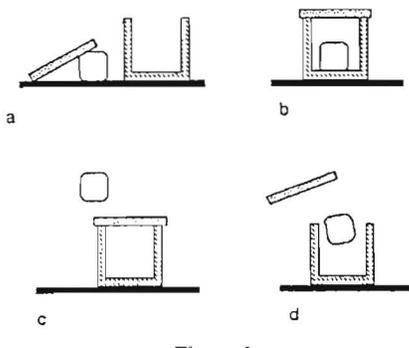


Figure 7 - The cube-into-the-box problem

Note that in this (planar) case the APF has four dimensions: the x and y coordinates of both moving parts

(the cube and the cover).

Monitoring the energy evolution during the action is a source of information for the mental simulation that can be used for analogical planning (for other related approaches to analogical planning see also Steels 1988, and Gardin and Meltzer 1989). A first way of exploiting this kind of information has already been described, and consists of generating negative icons (typically corresponding to local minima) during analogical reasoning (figure 6).

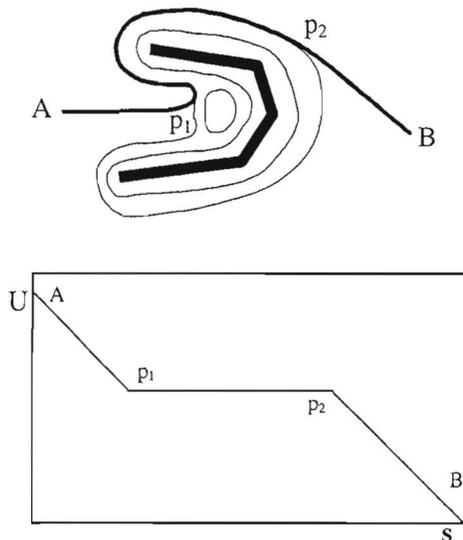


Figure 8 – Avoiding local minima by following equipotential lines.

An alternative way to exploit the potential energy evolution is the method of avoiding local minima by following an equipotential line of the APF. The field has necessarily closed equipotential lines surrounding obstacles. Even “concave obstacles” (typical sources of local minima) can be avoided by following one of the equipotential lines around them (figure 8, upper). It can be shown that a path for getting over an obstacle may always be reduced to an approach path (decreasing U), a leaving path (decreasing U), with an equipotential path in the middle (steady U) while getting around the obstacle without climbing the field (figure 8, lower). When the modelled state of affairs is complex, the space is multi-dimensional, but this principle still holds. Therefore, in many cases in which multiple objects are involved, deadlocks can be heuristically avoided by choosing a multi-dimensional path which follows an equipotential line of the global field. Consider the case of two mobile robots trying to cross a narrow corridor in opposite directions (figure 9). Robots repulse each other, have similar “strength” but casual speed and initial state, so that there is no symmetry in the model. Like in then example of figure 7, the APF has four dimensions. A free relaxation on the APF would give rise to a deadlock inside the corridor; if

we monitor the energy and change the behaviour of the two robots in order to keep constant the energy (i.e., choosing an equipotential line) inside the corridor (fig. 9, lower) we succeed in the task without deadlocks (fig. 9, upper). Unfortunately, when the APF has dimensions greater than 2 there are, normally, infinite equipotential trajectories. Nevertheless, if we are not interested in an *optimal* trajectory to reach the goal, this fact is not a serious drawback.

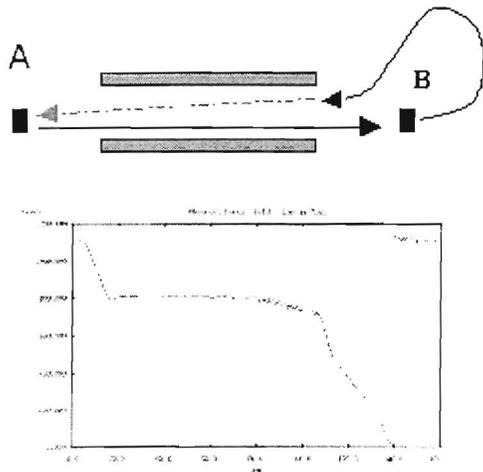


Figure 9 – A four-dimensions planning problem

Summing up, we have discussed so far three strategies for APF based diagrammatic reasoning about actions:

- using a priori schemata based on attracting and repulsing icons (fig. 5);
- iterative generation of repulsive icons (figs. 6,7);
- following equipotential lines (figs. 8,9).

APFs in Perceptual Tasks

In the introduction we said that image-like representations can help in relating perception with reasoning. Until now we have considered reasoning tasks performed on representational structures (icons) that could be imagined to be directly linked to the outputs of some perceptual module. Let us consider now the more ambitious case of reasoning about *real* scenes, i.e. the case in which icons come from (or are grounded on) perceptual data. We shall consider some examples, in which the energetic metaphor adopted could be used in perceptual tasks, in order to show the interrelation between perception and reasoning in this kind of models.

In the following subsections, we present two examples. The former concerns a constraint satisfaction problem in which APFs are used in low level perception, in order to deal with uncertain perceptual data (“uncertain icons”) coming from sensors. The latter is a higher level example, in which perception of dynamic scenes is interleaved with

the diagrammatic reasoning techniques described previously.

APFs Deal with Uncertain Perceptual Data

A significant role for energy relaxation in the mental simulation is that of finding an equilibrium point between “uncertain icons”, or, better, multiple fragments of a single icon affected by uncertainties. This is relevant when a certain diagram cannot be considered perfectly consistent with reality. APFs are natural interpolators and are able to find solutions which satisfy multiple uncertain constraints. Consider the problem of localising a robot *R* basing on angular measurements. A certain number of reference landmarks (natural or artificial) b_1, b_2, \dots are placed somewhere on walls, at known positions. The robot measures the angles $\alpha, \beta, \gamma, \delta, \dots$ (figure 10, above).

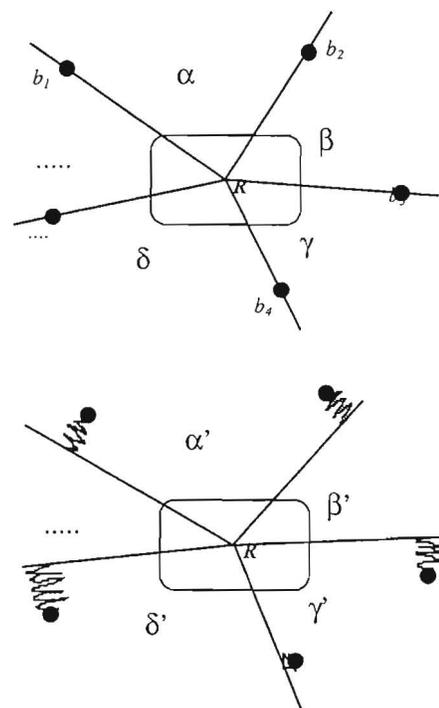


Figure 10 - Localisation by means of energy relaxation

The trigonometric algorithm for determining the position and orientation of *R* starting from measures of only *three errorless* angles is straightforward. In real cases, available measurements are redundant and affected by errors; moreover, the absolute position of reference landmarks is not exactly known as well. It is known that, due to the non linearity of the problem, little instrumental errors (e.g., 0.1°) may cause localisation errors of tens of centimetres or more, even in small environments. Finding a good estimate of the localisation of *R* in presence of redundant measures is a very complex task, which is usually faced either with non-linear mathematical programming, or with non-linear minimisation, or using Kalman filtering (Dean and

Wellman, 1991).

Using the iconic approach, we may think that all the “rays” joining R to landmarks form an icon (fig. 10). In a totally known environment, each ray pierces its own landmark (fig. 10, above). In uncertain environments, each landmark can be considered as an attracting entity for the corresponding ray (fig. 10, below). By letting the model “to relax”, R finds an equilibrium point at a global minimum, corresponding to an estimate of the position. With a fairly accurate localisation, the complexity of this task is so small that can be easily implemented to work on-line (it is currently part of the autonomous robot developed at DIST - Giuffrida *et al.* 1996).

Diagrammatic Reasoning for Perceptual Tasks

The APF metaphor can be extended to more complex, model-based perceptual processes, even of dynamic nature. In the localisation problem described above, the compositional properties of the APF are used “statically” (no “real” time is involved, the evolution of the model is a sort of off-line reasoning carried out “on the fly” as soon as perceptual data are available). In more complex perceptual processes, this mechanism can be mixed with the generative behaviour of the APFs discussed above. Starting from a series of partial views (partial snapshots, incomplete icons, like, for example, a 2-D projection of a 3-D scene) not containing the whole world description, it is possible to construct an icon-based description of the whole evolving scene.

The idea is that such internal model (a set of n -dimensional icons) is “attracted” by these partial/incomplete icons, as soon as they become available from perception. Since these last icons are, in general, not completely consistent (due to measurement errors, e.g. camera distortions, or because they are not exactly synchronised, etc.), the relaxation method described in the previous section is first used. Then, the internal model evolves toward the next set of incomplete icons. To do so, *a priori* knowledge about the model is used, e. g. using negative and positive icons, to avoid deadlocks, to regularise and interpolate data, to satisfy constraints, and to *complete perceptual data* (e. g., solving ambiguities due to lack of perception).

Suppose we want to reconstruct the real-time movement of a humanoid figure basing on *markers* put in his/her joints, seen from two lateral views. In this case, at any instant, the source of information is alternatively redundant and insufficient: usually, every point has more measurements than necessary, but sometimes a marker can become invisible from one of the points of view (figure 11).

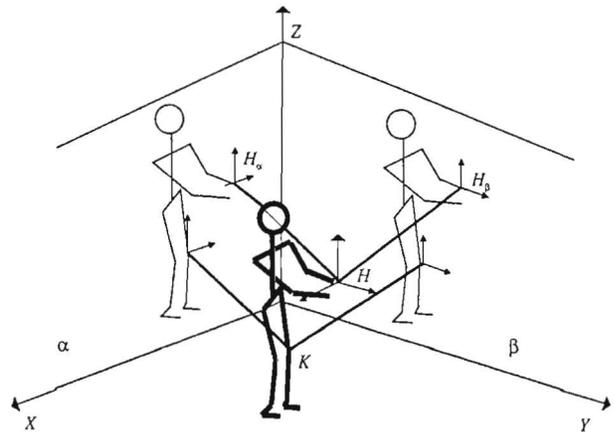


Figure 11 - Reconstruction of a complex movement

At any relevant time instant, the two lateral views are (incomplete) icons that can be metaphorically enhanced with “rays” perpendicular to the α and β planes starting from the projections of each marker (H and K in fig. 11). The intersection of each couple of “rays” should correspond in the model to the position of the corresponding marker in the world. In reality this does not happen, due to measurement errors. However, it is always possible to consider the representations of each marker in the model as “attracted” by its couple of “rays”. Moreover, each joint has the constraint of being connected to another joint by a limb of fixed length, so that the attractive affect to each joint propagate their effects to the whole kinematic chain.

The mechanism above is a relaxation process activated at every new snapshot, in order to give an *estimate* of the actual posture. Actual postures become intermediate, complete icons of the movement (key frames). The evolution of the model towards the next global icon is demanded to the relaxation process in the APF, acting as interpolator. However, this last process can benefit from the iconic/APF metaphor since it is possible to enrich the representation with *a priori* “analogical” knowledge for making the interpolation easier. Consider for example the case in which only one lateral view is available, so that a certain movement has an ambiguity. It is possible to represent a set of constraints (e. g., limit angles in the rotation of some joints) as negative icons which disambiguate the reconstruction by “taking away” the humanoid from inconsistent postures. For the same purpose it is possible to introduce concepts like *inertia* or *persistence*, and so on.

The mathematical complexity of such a model is sufficiently tractable so that it is possible to reconstruct the description of movements of a humanoid like in fig. 11 in real time using a DSP. The series of icons generated in this way could be the basis for a higher level process able to recognise and classify the movement. For sake of brevity, we omit the mathematics involved (kinematic

chains). It is worth noting that, besides the classical mechanics approach to the problem, alternative methods based on energy metaphors can be used (Mussa Ivaldi *et al.* 1988).

Conclusions

In this paper we have proposed a framework for reasoning about actions. This framework is diagrammatic and intrinsically dynamic, in the sense of being based on dynamic internal simulations analogous to a "mental cartoon". We claim that this approach can help to face in a natural way many spatial reasoning problems. The representation adopted is not universal: it makes easy to deal with certain kinds of spatial knowledge, but not with others. However, we maintain that this kind of model can be integrated with more traditional symbolic approaches, in order to build hybrid reasoning systems (Ardizzone *et al.* 1993). The emergence of discrete singularities from the evolution of the model (e.g. local minima in APFs) can help in associating symbols to the mental simulation. In this respect, this framework could be considered as an intermediate level of representation connecting perception and control on the one side with higher level propositional reasoning on the other side.

Acknowledgements

The authors would like to thank Prof. Piero Morasso for the stimulating discussions about many of the ideas underlying this paper, and Dr. Maurizio Piaggio for fruitful discussions.

References

- Ardizzone, E., A. Camurri, M. Frixione, R. Zaccaria 1993. A Hybrid Scheme for Action Representation. *International Journal of Intelligent Systems* 8(3): 371-403.
- Arkin, R.C. 1990. Integrating Behavioral, Perceptual and World Knowledge in Reactive Navigation. *Int. J. Robotics and Autonomous Syst.* 6 (1-2): 105-122.
- Bosch Hausgeräte GmbH 1990. *SMS 7062 Dishwasher Operating Instructions*. München: Robert Bosch Hausgeräte GmbH.
- Block, N. ed. 1981. *Imagery*. Cambridge, Mass.: MIT Press.
- Chandrasekaran, B., H. Simon eds. 1992. *Reasoning with Diagrammatic Representations*. Papers from the 1992 AAAI Spring Symposium, Menlo Park, Calif.: AAAI Press.
- Clerici, G. 1972. *Il tennis facile*. Milano: Mondadori.
- Connolly, C.I., J.B. Burns, R. Weiss 1990. Path Planning using Laplace's Equation. *Proceedings of IEEE Conf. on Robotics and Automation, ICRA '90*: 2102-2106.
- Dean, T.L., M.P. Wellman 1991. *Planning and Control*, San Mateo, Calif.: Morgan Kaufmann Publ.
- Fauconnier, G. 1985. *Mental Spaces*. Cambridge, Mass.: MIT Press.
- Feldman, A.G. 1966. Functional Tuning of the Nervous System with Control of Movement or Maintenance of a Steady Posture. II. Controllable Parameters of the Muscle. *Biophysics*, 11: 565-578.
- Forbus, K. 1995. Qualitative Spatial Reasoning: Framework and Frontiers. In Glasgow *et al.* eds. 1995.
- Funt, B.V. 1980. Problem-Solving with Diagrammatic Representations. *Artificial Intelligence* 13(3): 201-230.
- Gardin, F., B. Meltzer 1989. Analogical Representations of Naive Physics. *Artificial Intelligence* 38: 139-159.
- Gelerntner, H. 1959. Realization of a Geometry-Theorem Proving Machine. In E.A. Feigenbaum and J. Feldman eds., *Computers and Thought*, New York: McGraw-Hill.
- Giuffrida, F., P. Morasso, G. Vercelli, R. Zaccaria 1996. Integration of Active Localization Systems in Vehicles Control for Real-Time Trajectory Tracking. *Proc. 1996 IEEE/SICE/RSJ Int. Conf. On Multisensor Fusion and Integration for Intelligent Systems*: 549-556.
- Glasgow, J., N.H. Narayanan, B. Chandrasekaran eds. 1995. *Diagrammatic Reasoning*. Cambridge, Mass.: MIT Press.
- Johnson-Laird, P. 1983. *Mental Models*. Cambridge: Cambridge University Press.
- Khatib, O. 1986. Real Time Obstacle Avoidance for Manipulators and Mobile Robots. *Int. J. Robotics Res.* 5(1): 90-99.
- Koditschek, D.E., 1987. Exact Robot Navigation by means of Potential Functions: Some Topological Considerations, *Proc. of the IEEE Int. Conf. of Robotics and Automation*, Raleigh, NC: 1-6.
- Koren, Y., J. Borenstein 1991. Potential Field Methods and their Inherent Limitations for Mobile Robot Navigation. *Proc. IEEE Int. Conf. on Robotics and Automation*: 1398-1404.
- Kosslyn, S.M. 1980. *Images and Mind*. Cambridge, Mass.: Harvard University Press.
- Kulpa, Z. 1994. Diagrammatic Representation and Reasoning. *Machine Graphics & Vision* 3(1/2): 77-103.
- Latombe, J.-C. 1991. *Robot Motion Planning*. Norwell, Mass.: Kluwer.
- Levesque, H. 1988. Logic and the Complexity of Reasoning. *Journal of Philosophical Logic* 17: 355-389.
- Loeff, A.H. Soni 1975. An Algorithm for Computer Guidance of a Manipulator in between Obstacles. *Transactions of ASME, Journal of Engineering for*

- Industry* 3: 836—842.
- Lozano-Pérez, T., and Wesley, M. A., 1979. An Algorithm for Planning Collision-Free Paths among Polyhedral Obstacles. *Communications of the ACM* 22(10): 560-570.
- Morasso, P., V. Sanguineti, G. Spada 1997. A Computational Theory of Targeting Movements based on Force Fields and Topology Representing Networks. *Neurocomputing* 15: 411-434.
- Mussa Ivaldi, F.A., P. Morasso, R. Zaccaria 1988. Kinematic Networks. A Distributed Model for Representing and Regularizing Motor Redundancy. *Biological Cybernetics* 60: 1-16.
- Myers, K., K. Konolige 1995. Reasoning with Analogical Representations. In Glasgow *et al.* eds. 1995.
- Slack, M.G. 1993. Navigation Templates: Mediating Qualitative Guidance and Quantitative Control in Mobile Robots. *IEEE Transactions on Systems, Man, and Cybernetics* 23(2).
- Sloman, A. 1995. Musings on the Roles of Logical and non-Logical Representations in Intelligence. In Glasgow *et al.* eds. 1995.
- Steels, L. 1988. Steps Towards Common Sense. *Proc. 8th European Conference on Artificial Intelligence ECAI-88*: 49-54.
- Tilove, R.B. 1990. Local Obstacle Avoidance for Mobile Robots based on the Method of Artificial Potentials. *Proc. IEEE Conf. on Robotics and Automation*: 566-571.