A Theory-Based Dynamical Model of Innovation Processes

David Lane Roberto Serra Marco Villani
Luca Ansaloni

Department of Social, Cognitive and Quantitative Sciences, Modena and Reggio Emilia University, Reggio Emilia, Italy

Key Words

Multi-agent-based model • Innovation processes • Economy • Complex artificial systems, design

Abstract

We present an agent-based model of innovation processes, based upon a theory of innovation by Lane and Maxfield. The theory inspires and constrains the features of the model, thus reducing the *embarras de richesse* that is one of the major methodological problems of agent-based modelling. Artefacts are produced by agents using recipes; the basic dynamics, absent innovation, is one of production and sales, where the external world supplies 'raw materials' and external demand. Depending upon the initial conditions, self-sustaining cycles of production and exchange can emerge among the agents. Innovation – that is, the generation of new recipes, in particular desired directions, called 'goals' – results in substantial modification of the system dynamics. Two innovation regimes are introduced: a 'lonely' mode, in which each agent tries to introduce new products by itself, and a 'relational' mode, in which two agents can improve their reciprocal knowledge and can decide to try to jointly develop a new artefact.

Simplexus

What is innovation? What factors determine whether we innovate or stagnate, and is it possible to predict future innovations?

These questions have vexed researchers studying socio-politico-economic structures in the information society. Lane, Serra, Villani, and Ansaloni point out in the current paper how others have attempted to model how innovations occur, how they develop, and their impact on society by using so-called agent-based models, ABMs.

In an ABM, an agent, or individual, is defined as an entity, plants and animals in ecosystems, vehicles in traffic, or the virtual agents of software and the computer gaming world, for instance. An ABM attempts to reconstruct, or simulate, the wider, or global consequences of the local activities and interactions of the individual agents. It usually consists of a framework within which the agents can interact and depending on their actions cause change either to agents or the framework itself. For instance, the interactions of plants and animals in a jungle ecosystem might be modelled so that the effects of population changes of particular species can be predicted.

In the case of innovation, the agents are the ideas of people, inventors and developers, the marketing people, advertisers, sellers and distributors, and the end-users. ABMs have advantages over models that attempt to predict global outcomes on the basis of population averages.

According to Serra and colleagues, however, ABMs often result in an *'embarras de richesse'*. The underlying problem is that an ABM allows the modeller too much freedom, they can enter almost limitless variables and tweak parameters endlessly. Moreover, such freedoms require a larger input data set to provide any meaningful output, and this is not always readily available. What is needed is a way to realistically constrain the model.

Copyright © 2005 S. Karger AG, Basel

Fax +41 61 306 12 34
E-Mail karger@karger.ch
www.karger.com

KARGER

© 2005 S. Karger AG, Basel
1424-8492/05/0024-0177
\$22.00/0
Accessible online at:
www.karger.com/cpu

Prof. Roberto Serra, Department of Social, Cognitive and Quantitative Sciences, Faculty of Communications and Economics, Modena and Reggio Emilia University, via Giglioli Valle 9, IT-42100 Reggio Emilia (Italy), Tel. +39 0522 522433
Fax +39 0522 522512, E-Mail serra.roberto@unimore.it

1 Introduction

Modelling social and economic phenomena is a major challenge. In recent years agent-based models (ABMs) have attracted considerable attention, as they allow one to escape from the constraints of the 'representative agent' and to account for agent heterogeneity [see e.g. 1–6 and further references quoted therein]. ABMs are well suited to bridge the gap between hypotheses concerning the microscopic behaviour of individual agents and the emergence of collective phenomena in a population composed of many interacting agents. Moreover, by requiring that hypotheses be stated precisely, ABMs may improve the development of the theory itself.

Attractive as ABMs may be, they still face formidable problems. In particular they allow the modeller too much freedom with respect to the design of agents and their interactions. ABMs often have too many variables and parameters. Moreover, while in some cases experimental data for model testing are available, in many interesting cases data are scarce and do not sufficiently constrain the model. In these cases, what is needed to get meaningful results is a principle which limits the degrees of freedom of the modeller. While Ockham's razor is of the outmost importance, it does not seem to be sufficient by itself for the purpose of limiting the *embarasse de richesse* of ABMs.

In this work we consider another approach, namely that of relying upon an existing theory of the phenomenon in order to constrain our modelling options. The theory, which is qualitative in nature, provides the basic entities of the model and predicates about their relationships. The model represents a simplified universe inhabited by (some of) the theoretical entities. Simplification is necessary in order to deal with manageable systems: we do not look for an all-encompassing model, but we rather imagine a family of models which capture different features of the theory.

The phenomenon which is considered in this paper is that of innovation, and the theory we refer to has been initially developed by Lane and Maxfield [7, 8], and is under further development within the UE-FET project Iscom (Information Society as a Complex System), whose approach is based on the interaction between theory development, case studies and model (see the project website www.iscom.unimo.it). We will refer for brevity to this theory as the Lane-Maxfield (LM) theory, while the model will be called I_2M (Iscom Innovation Model).

A very brief outline of some of the major features of the LM theory will be presented in section 2. The theory constrains the model, and the main constraints will be summarized at the end of that section. Section 3 describes the main features of the model. Section 4 describes the characteristics of a software simulator that implements parts of the model, while section 5 presents the results of some simulations. Finally, section 6 contains comments on the results achieved so far and indications for further work.

2 Indications of the Theory

In the work of Lane and Maxfield [7, 8], artefacts are given *meanings* by the agents that interact with them, but these meanings cannot be understood without taking into account the *roles* which different agents can play. Thus, artefacts may be given different meanings by different agents, or by the same agents at different times.

Therefore, LM theory can be seen as a theory of the interpretation of innovation. According to LM, a new interpretation of a potential artefact functionality can be put forth in the context of so-called generative relationships. By interacting, a few agents come to invent and share this interpretation, based on the discovery of different perspectives and use of existing artefacts. The generative potential of a relationship may be assessed in terms of the following criteria:

Fortunately, there is an alternative that retains many of an ABM's advantages but avoids data embarrassment, by reducing the options available to the modeller. The method involves applying an existing theory of a phenomenon to constrain the ABM. The theory provides the basic characteristics of the agents and limits the variables associated with how they interact. In other words, it simplifies the framework so that wildly unrealistic or abhorrent agents and behaviour are not included in the model. The approach was developed initially by Lane and Maxfield and is currently being developed further within the UE-FET project Iscom (Information Society as a Complex System). Fundamentally, it is based on the interaction between theory development, case studies, and the model.

The researchers have now applied the Lane-Maxfield (LM) theory to the world of innovation. The key to making LM work in this context involves contextual semantics – meaning is attached to functional artefacts, the innovations themselves, by the agents, and this meaning can be different depending on the characteristics of the agents and the time at which they consider an artefact. An artefact's meaning is essentially what agents do with them, and the agents' various roles are defined in terms of their artefacts. For instance, the meaning of a cellphone varies depending on whether the agent handling such an innovation is the designer, salesperson, or the person making their first call. At its simplest, one might envisage a chain of interactions between agents. The novel cellphone is designed by one agent, who passes on the design to the next agent, the manufacturer, who then distributes it to the salespeople, who sell it to final agent in the chain – the user.

However, it is obvious that such a simplistic chain ignores dozens of agents that might have inputs along the chain, or be running parallel processes, such as designing the packaging or copywriting an advertising campaign. Manufacturers can also

- heterogeneity: the agents are different from each other, they have different features and different goals; the heterogeneity is not so intense as to prevent communication and interaction
 - aligned directedness: the agents are all interested in operating in the same region (or in neighbouring regions) of agent-artefact space
 - mutual directedness: the agents should be interested in interacting with each other
- Moreover, Lane and Maxfield discuss two further features that deal with organizational issues, i.e. permissions and action opportunities. The former refers to the fact that the agents are authorized to engage in such a relationship, the latter to the possibility of moving from 'talk' to action.

Lane and Maxfield further argue that, in a situation where innovations happen at a very fast pace, predicting the future is impossible; so a better strategy would be to identify those relationships that have the potential for generativeness, and to foster them in order to effectively explore the new opportunities they can give rise to. It is therefore very important to be able to estimate the 'generative potential' of the existing and prospective relationships.

The LM theory of innovation is highly sophisticated in describing the interactions between different players in innovation processes, and it cannot be entirely mapped onto a specific computer-based model. Therefore, the modelling activity aims at developing models that are based on abstraction of some key aspects of this theory, which is of a qualitative nature.

The basic requirements for the model derived from the theory can be summarized as follows:

(1) The meanings of artefacts must be generated within the model itself: since the LM theory claims that new meanings are generated through interactions among agents and artefacts, it would be inappropriate here to resort to an external oracle to decide a priori which meanings are better than others.

(2) The roles of agents must also be generated within the model: indeed the LM theory claims that new roles are also generated through interactions among agents and artefacts.

(3) Agents must interact with artefacts and with other agents: interacting with artefacts only would prevent the possibility of describing agent-agent relationships.

(4) An agent should be able to choose the other agents with whom to start a relationship; in general, an agent will be able to handle a finite number of relationships at a time, and it will choose a subset of the other agents as its partners. Agents must be allowed to cut a disappointing or unsatisfactory relationship to look for a better one.

(5) An agent should be able to have different degrees of interaction with another agent: in this way it will be possible to tune the intensity of these relationships.

(6) Some agent-agent relationships should be generative in character, i.e. they should be able to lead to new attributions (for agents and/or artefacts) and should respect the criteria for generative potential described above.

3 Model Description

We introduce a model where innovation happens in an 'artificial world' [9], which is inhabited by agents which 'produce' artefacts, which in turn can be used by other agents to build their own artefacts. An agent can produce several artefacts for different agents (and it can sell one type of artefact to several different customers).

The meaning of artefacts is just what agents do with them, while the role of agents is defined by which artefacts they produce, with whom, and for whom. The role of agents is also partly defined by the social network they are embedded into. In order to describe these features, and to have them generated in the model, without being enforced a priori, we envisage a network of agents that are producers of artefacts, which are in turn used by other agents.

collaborate, or employees move within an industry. Additionally, there are further links beyond the end-user associated with the telecommunications provider and even the disposal or recycling agency when the cellphone reaches the end of its useful life. There are also dozens of agents peripheral to the chain, company shareholders, rival companies, and users of other types of cellphone.

Understanding who is who within a model world of innovation, and that a designer designs, or a seller sells, for instance, is only the first step to creating a model. If a simulation is to be useful, it has to be digitized, which means creating an artefact space and populating it with mathematical entities analogous to the real agents of innovation. How the agents interact and the functions they perform can in turn be represented by algebraic operations. The operators are equivalent to technology, so that division might represent an innovation in a world of addition and subtraction. Such an operative innovation might represent the creation of cellphones with an in-built digital camera, for instance.

Serra and his colleagues have formalized this notion of innovation and artefact space using the principles of ABMs and constraining their model using the LM theory and describe a prototype simulator. This comprises a computational toolkit that creates the initial network of agents and legacy artefacts, a simulation engine, and a module to visualize the main simulation results. Using this setup, they can simulate a wide range of situations and conditions in which agents interact to build innovative or novel artefacts on a single personal computer, or for more complex simulations, a cluster of PCs.

The researchers explain that in order to innovate and so introduce new products, agents must have goals and develop the recipes to produce them. They have tested two different ways an agent might find a new goal – 'imitation', which as the name implies is copying someone else's idea, re-

So the ‘strong ties’ between two agents are mediated by a chain of artefacts: for example, agent A1 produces artefact O1 which is used e.g. by agent A2 to produce artefacts O2 and O3, and by agent A3 to produce artefact O4. O2, O3 and O4 are used by other agents. If we describe the agent network, then there is a strong (directed) link from A to B if A produces an artefact which is used by B: the network of agents may well have loops, and an agent may have multiple outgoing and multiple incoming links.

There are also weak ties between two agents (‘acquaintances’) which refer to the fact that agent A knows something about agent B (e.g. its products, or part of its providers): so the social neighbours of a given agent are a superset of its providers and customers.

A key point is the structure of artefact space. For modelling purposes, we have considered different alternatives: binary coding as it is done in classifier systems, or λ -calculus as in the Alchemy model [10, 11] or numbers (either natural or real). What is required is that the space has an algebraic structure, and that suitable constructors can be defined to build new artefacts by combining existing ones. We concentrated on the real number representation and the use of mathematical operators because it is more compact than the binary representation and simpler than the predicate calculus and the λ -calculus. A full comparative evaluation of the features of the different representations lies beyond the purpose of this paper.

Therefore, the agents are ‘producers’ of numbers by combination of other numbers. The mathematical operators which are available to an agent represent its ‘technology’, so a major technology shift might be described by the appearance of new operators, like e.g. the discovery of division in a world where addition and subtraction were the only available technologies.

It is well known that technological constraints limit the set of possible artefacts, a

feature which is sometimes described by defining a fitness landscape on the artefact space; however, in our model the value of an artefact must be related to what other agents do, so assigning a priori a given fitness would not be a suitable alternative.

A more precise description of the model is given below.

3.1 Entities' Description Artefacts

The artefact Ar_i ($i \in [1, N]$) is described by:

- the identifier $i \in \mathbb{N}$, which indicates in a unique manner the artefact itself
- the integer value $name \in \mathbb{N}$, which identifies the artefact type and is the main value manipulated by the agent recipes
- the identifier of the agent that produced it
- the level of its stock
- the identifiers of the agents that acquire it

We foresee the possibility to develop in the future a richer representation of the agent type, by resorting to a vector of integers (or real numbers) instead of a single scalar quantity.

Agents

The agent Ag_j ($j \in [1, M]$) is described by:

- the identifier $j \in \mathbb{N}$, which indicates in a unique manner the agent itself
- the number of recipes $n_{recipes} \in \mathbb{N}$ owned by the agent
- the number of goals $n_{goal} \in \mathbb{N}$ and the list $list_{goal} \in \mathbb{N}^{n_{goal}}$ ‘owned’ by the agent
- the ‘catalogue’, i.e. a list of existing recipes, available to the agent
- the list of the agent acquaintances
- the list of the known artefacts
- the real vector $vote \in \mathbb{R}^M$, which measures the quality that the agent attributes to each relation (with other agents) it maintains
- the real value $activity \in \mathbb{R}$, which measures the ‘wealth’ of the agent

fashioning the standard cellphone in new colours would be considered imitation, for instance. The second goal-finding strategy is ‘imitation and jump’, which involves more than simply emulating another product, but also adding value, the cellphone with an integral digital camera, for example.

Imitation alone is unlikely to boost mathematically the agent’s value in the model, or in the real world increase sales. The researchers have run several simulations and present case studies rather than a statistical analysis in the current paper. Increased agent sophistication is reflected in the following:

- (1) Agents operate without innovation.
- (2) Agents perform innovations, either in suppliers or in goals and recipes, each agent knowing all the existing artefacts.
- (3) Agents perform innovations, either in suppliers or in goals and recipes, each agent knowing only the artefacts produced by a limited number of other agents.

In case 1, the agents do not perform any kind of innovation and never change their providers. Two different scenarios emerge from the model: without an external demand, all agents quickly tend to zero and never again rise above the baseline. With an external demand, each agent’s value plateaus and remains in this steady state indefinitely.

In case 2, agents are aware of competing artefacts, can imitate them, but after an initial flurry of activity all once again tend to zero (scenario without external demand). However, if they are allowed to change their providers, self-sustaining structures can emerge, their activity is sustained and the overall value of the agents plateaus but that of individuals fluctuates above and below this plateau indefinitely. Removing a single agent at this point, however, can disrupt the self-sustaining structures and introduce an avalanche of changes of different size.

In all the variations of the first two cases innovation comes to a halt either through the achievement of a stable steady state or

- a set of parameters that control the agent propensity to innovate, to enlarge its own acquaintances, to participate in joint projects, and so on (in the future, we will sometimes refer to this set of parameters as the agent ‘style’)

Recipes

The third object we define is the production recipe. The k -th recipe of the agent Ag_j , $R_{c_{jk}}$ [$k \in (1, Ag_j \cdot n_{recipes})$] is described by:

- the number $n_{input} \in \mathbb{N}$ of needed inputs
- the ordered list of identifiers of the input artefacts
- the ordered list of operators, whose number is equal to $n_{input} - 1$
- the identifier of the product
- the integer value $prod \in \mathbb{N}$, which states the number of tokens created at each recipe production
- the real value $cost \in \mathbb{R}$, which measures the cost required to produce the product associated with that recipe
- the time $t_{last} \in \mathbb{N}$ elapsed from the last time the recipe has been active, i.e. actually used for production

It is useful to define a further entity, the recipe support, which is simply a subset of recipe features:

- the number $n_{input} \in \mathbb{N}$ of needed inputs
- the ordered list of operators, whose number is equal to $n_{input} - 1$.

This concept allows the separation between the ‘fundamental’ structure of one recipe (i.e. the recipe support, which stores the construction modalities that are the basis of the recipe itself) and one of its instantiations (which leads from the particular input artefacts to the output one).

External Environment

The model presented here is not necessarily a closed one, even if, as we will see, it can behave in a self-sustaining mode, without need of external inputs. However, in order to allow a wider generality of the mod-

el, we will assume that there may be an ‘external world’, which represents other agents which are not included in the model (external customers) or artefacts which are not produced within the model (e.g. raw materials). We can call all these entities ‘environment’.

More specifically, we identify:

- a set of ‘raw materials’ (artefacts whose quantities are neither limited nor associated with some particular agent; they can be used as input to production recipes)
- a set of ‘external reward regions’, areas where artefacts are acquired by entities not belonging to the simulated market system

3.2 Model Dynamics

We first provide an overall description of the model behaviour. Agents transform artefacts; in agent-artefact space they are defined by their production recipes and by their relationships with other agents. Each agent can produce more than one product, and for each product there is at least one production recipe. For simplicity we state that at each simulation step only one recipe is active for a given product. Agents have also ‘goals’, that is particular names that the agent aims to produce. Furthermore, agents are able to participate in joint ventures, to enlarge their acquaintances, and to give birth to new agents.

A schematic sketch of the possible agents’ actions is given below.

At each time step (updating is asynchronous):

a randomly chosen agent is selected for updating

(a) for each recipe, the agent looks for the required inputs

(I) if the inputs are found (i.e. they have nonzero stock level), the recipes’ outcome is produced and inserted immediately into the stock, while the corresponding input stocks are decreased

(II) if the inputs are not found, the agent searches, among the artefacts known to it,

as all the agents reach value zero and cease trading. In case 3, however, where innovation can be accompanied by a leap of faith to achieve goals with added value, the picture is altogether more complicated. The jumps are simulated by adding a randomization factor into an agent’s goals to mimic the notion of a bright spark, a jump in the innovation process. Now, not only are agents imitating each other’s cellphones and connecting in various ways, they are developing more new features, such as making the in-built digital camera a video camera too, and adding video playback to the device. As such, certain agents in the network can out-perform the others and make value leaps as the simulation progresses.

The team has also investigated how these simulations behave when there is limited knowledge of rival agent’s artefacts, the effects of small fluctuations, and the effects of carefully choosing partners for innovation. Despite this seeming sophistication, the present model represents a very simplified view of social networks, based upon chance. One of the major future directions of research, say the researchers, would be to involve intentionality in the model and also to change the nature of the interactions between agents to mimic real-life business or other networks. They are now improving the computer code to accelerate the simulations, so allowing more agents to be considered.

Innovations themselves cannot be predicted easily. However, the Serra model can provide hints concerning the ‘signatures’ of innovations which may lead to large changes in agent-artefact space, thus allowing one to recognize at an early stage those phenomena which might have major consequences.

David Bradley of Sciencebase.com

another object with the same name, whose stock level is different from zero

(i) if this artefact is found, it substitutes the old one; return to step (1.a.I)

(ii) if this artefact is not found, pass to another recipe; in this case, a counter associated with the unused recipe is incremented (this counter will be used to cancel recipes which have not been used for several times)

(b) the following actions need not be performed at each time step. Depending upon the value of suitable parameters, and upon the level of its strength, the agent may perform one or more of the following (the decision about which ones is taken in a stochastic way)

(I) try to innovate by itself its catalogue by

(i) choosing a goal

(ii) trying to realize it, i.e. to find a recipe for the goal (or for a product close enough to the goal, see below section 3.2.1)

(iii) producing it, if a successful recipe is found

(II) try to innovate its catalogue together with another agent, by

(i) choosing a partner

(ii) with the partner, choosing a goal

(iii) with the partner, trying to realize it

(iv) sharing the obtained results, if the innovation process succeeds

(III) try to enlarge its acquaintances

(IV) cancel the less useful recipes

(c) the agents' activations are updated

Moreover, new agents can be created, according to one of the following rules:

(d) randomly, with a given probability

(e) in order to substitute the dead agents

(f) according to the internal dynamics

Now we present in detail the actions just sketched above.

3.2.1 Goals

In order to survive (i.e. to keep its strength above a given threshold) it is useful for an agent to find artefacts that could have high sales. Therefore, identifying a

proper goal (the choice of the artefact that has to be built) is crucial for the agents.

When its turn comes, an agent which is rich enough (that is, whose strength is high enough) can try to create a new recipe. Let us consider 'lonely' agents, which try to do this basically by themselves; the agent:

- looks at the portion of artefact space which it knows
- updates the set of its goals
- chooses a specific goal
- tries to come close to that goal by using the available operators on the available inputs and recipes
- if, using the available inputs, it comes close to the goal (within a given range) it puts the corresponding recipe in the set of active recipes

Heuristics for goal setting include:

- locate a region of artefact space which is populated by several artefacts (and that therefore is likely to be able to sustain a high number of them)
- explore an existing reward region looking for higher reward points
- explore new regions of artefact space, e.g. by mutating an existing artefact or by interpolating between some of them
- respond to the request of another agent
- examine another agent's recipes and find some input product which can become a goal for the first agent

3.2.2 Recipes

Agents transform artefacts by means of their production recipes. Each agent can produce more than one product, and for each product there is one production recipe. A recipe is a function $R_{c_{ij}}: In \rightarrow Out$, where $In = \{x_1, x_2, \dots, x_{n_{input}}\}$ is the vector of the names of the input artefacts and Out is the name of the product.

The total transformation is composed by a set of simpler transformations, called operators; the presence (or absence) of particular kinds of operators defines the 'technological level' achieved by the recipe. The operators are not limited to the basic numerical operators (addition, subtraction,

multiplication, division), but more generally are any sort of functions

$$Op: E \subseteq \mathbf{R}^m \rightarrow \mathbf{R}$$

like for example exponential, mean, maximum, or function restrictions (where the operand domain is restricted to certain intervals, e.g. addition with names $\in [0, 120]$, multiplication between artefacts with names ≤ 10 , etc.). If we assume that $Op = \{op_1, op_2, \dots, op_{n_{input}-1}\}$ is the list of the operators, the result of a recipe could be written as:

$$Out = ((((((x_1 \cdot op_1 \cdot x_2) \cdot op_2 \cdot x_3) \cdot op_3 \cdot x_4) \cdot op_4 \cdot x_5) \dots) \cdot op_{n_{input}-1} \cdot x_{n_{input}}))$$

Note that, in the present notation, the order of the application of operators is from op_1 to $op_{n_{input}-1}$. Changes of the order of the input artefacts, as well as changes of the operator's order, generally lead to changes in the recipe output. Of course, changes of the names of the input artefacts also lead to changes in the recipe output. All these alterations nevertheless could be done without modifications of the recipe support, which remains the fundamental structure of the recipe.

Agents are capable of changing the recipe repertoire by:

(1) eliminating one recipe and the corresponding product (at no cost)

(2) introducing a new recipe, by means of:

(a) the change of the providers (new recipe, same recipe support, same input names)

(b) the change of the order of the inputs (new recipe, same recipe support)

(c) the change of the order of the operators (new recipe, new recipe support)

(d) modification of the choice of operators and possibly of their number (different recipe, different recipe support).

It is possible to introduce new recipes with at least two different approaches:

- acting alone and using only one's own recipes

- initiating a collaboration with another agent and combining the knowledge of both partners

A recipe may be active as long as it has input artefacts; of course, if one or more of its inputs are absent from the system, the recipe does not operate and its output is not produced. A recipe whose product does not have any client for a certain number of time steps is removed and the product disappears.

When a new recipe is first introduced, the corresponding agent inserts its product in its catalogue, making it known to its acquaintances. If at least one customer is found the recipe is activated, and remains active as long as it has at least one customer. The first time no customers are found the recipe does not act, although it remains in the accessible catalogue for some time t_{forg} . Recipes that have been constantly inactive for a certain number of steps are removed ('forgotten').

Recipes which are regarded as useless (e.g. the recipes that do not make a sensible contribution to the agent activity) or even damaging (those that make a negative contribution) may be removed by the owner, so that details (e.g. parameters) depend upon the agent 'style'. Moreover, an agent could choose to eliminate the less performing recipes to leave room for more performing ones.

Note also that, in the model described so far, as well as in the simulations shown in section 5, artefact space is 'flat', i.e. uniform, although some regions may be unreachable at a certain time because of a lack of proper inputs or operators. The model could however also deal in a straightforward way with intrinsic difficulties in reaching some regions of artefact space ('technical barriers'), by modulating the success probability of an innovation using some suitable external function $F(x)$ (e.g. $|\sin x|$): when an innovation is tried, the probability of success is proportional to $F(x)$.

3.2.3 Joint Development of New Products

In order to align and coordinate their actions, two or more agents, which have goals which are close in artefact space, can cooperate; if one agent alone is not able to reach the interesting region, the combined efforts of two agents could be successful.

A way to combine the efforts is that of sharing the agents' knowledge (input artefacts, recipes). The agents could create a new recipe (for example, combining their recipes by means of genetic operators) which is able to reach the interesting region; after this action, they could divide the obtained recipe into two different recipes in such a way that the product of the first recipe is one of the input artefacts of the second recipe. Finally the first agent becomes the owner of the first recipe and the second agent becomes the owner of the second recipe (one recipe for each agent). In this way they are able to share the success of the research results.

3.2.4 Artefacts' Stock Level

Artefacts are characterized, among other variables, also by their stock levels, that is the 'quantity' of the available tokens. When a recipe uses one artefact the corresponding level decreases, whereas when the appropriate recipe produces the artefact, the corresponding level increases.

Generally speaking, a recipe combines a certain amount of each input product to generate a certain amount of output. A precise quantitative description of these aspects would lead to further complications. Furthermore, we expect that this complication is related more to the specific aspects of the production chain than to the model aims, i.e. the study of innovation processes. On the other hand, we feel it necessary to maintain the fundamental distinction between types (i.e. names of the artefacts) and tokens (i.e. quantities). This is an example of how the theory constrains the model! In order to meet this need while keeping the complications at a minimum, we consider

recipes where the quantity of each input is one, while the quantity of the output product is a property of the recipe itself (described by the variable '*prod*').

If the number of produced tokens is equal to the number of inputs, the total number of tokens is conserved; otherwise, this number is allowed to change. Therefore, if the value of this variable is greater than the number of inputs the total number of artefacts (tokens, not types) will tend to increase in time.

The agents control the stock levels: their aim is to avoid stock depletion and to avoid a useless surplus. There are several methods to obtain these aims:

- to allow production only when the stock level is equal to zero
- to allow production only when the stock level is lower than a given value (e.g. *prod*)
- to allow production depending on the artefact sales (e.g. by comparing sales trend with stock levels)

3.2.5 Activation Dynamics and Agents' Wealth

A real variable, called activation, is associated with each agent, and measures its 'wealth'. At each time step, the agent activation is updated; it receives positive contributions from those products that are sold, while it has to 'pay' for all its recipes. Whenever an agent performs a change (e.g. by introducing a new product or changing a recipe) it has to pay some activation in order to be allowed to try. A stochastic mechanism then determines whether the attempted modification happens or not: this simulates the effort that can be put in research and development to discover new recipes.

The formula that describes activation dynamics is:

$$Ag_j.act(t+1) = Ag_j.act(t) + \sum_{k \in \text{sold.products}} n_{jk}(t) - Ag_j.act(t) \cdot \sum_{i \in \text{active.recipes}} Rc_{ji}.prod \cdot Rc_{ji}.cost - \lambda Ag_j.act(t) \quad (1)$$

where $n_{jk}(t)$ indicates the number of sales of product k (belonging to agent j) at time t . The final term is necessary to make the activation of an agent which never acts vanish.

Of course, different formulae for the activation dynamics can be tested; simulation results will provide indications about the better ones.

Agents whose activation is equal to or smaller than zero are removed from the system. New agents could be generated according to rules whose definition depends upon the aims of the simulations. For example, some suitable rules are:

- to substitute the removed agents (in order to maintain the agents' number)
- to introduce new agents at each step according to a given probability distribution (in order to simulate the appearance of new agents coming from outside the system)
- to introduce new agents according to internal dynamics, for example depending upon the total system wealth, or the wealth of a particular agent that can generate some spin-off

3.2.6 Relationships among Agents

Agents can use their knowledge about agent-artefact space, which is in general limited (an agent may not know all the existing artefacts, all the existing agents, all the goals and the recipes of the known agents).

An agent can also generate new goals or remove some old ones, with the aim to increase its activation. The definition of a new goal can be based upon the knowledge of the other agents' requirements and of the existing products, and upon the requirements of the external world.

An agent is also endowed with a 'social network'. There are two different kinds of interactions among agents: (1) relationships with providers and customers/clients, mediated by artefacts, and (2) relationships with the agent's acquaintances

('weak ties'), in which some exchanges of information can take place.

Each agent has its own measure of the quality of its relationship with other agents, which increases with the frequency of the positive events that take place during their interactions (purchases, sales, exchange of information), while it decreases with the frequency of negative events (wrong information, mistakes, poor activity income, failure in joint developments).

A reciprocal high vote could allow two agents to:

- exchange information (acquaintances, goals, operators,...)
 - share (part of) their recipes
 - create alliances in order to reach interesting zones of agent-artefact space
- Agents can invest part of their activity, in order to:
- enlarge their acquaintances
 - look for new providers
 - improve their knowledge of some agent with whom they hold a high-ranked relationship
 - increase the number of artefacts they know

To improve their relationships, an agent can inform other agents about (some of) its desired inputs. If the interaction has been successful, the agent lets its providers 'know more' about itself – e.g. other desired inputs or its production recipe(s). Moreover, reciprocity leads providers to inform their customers about themselves.

4 The Prototype Simulator

In order to test the behaviour of the model and the suitability of some mechanisms that have been proposed, we developed a preliminary, simplified version of the simulation environment, which has three main modules:

- a specialized initial condition builder, whose aim is to create the initial network of agents and artefacts
- the simulation engine
- a tool to visualize the main simulation results

The interaction among the modules allows one to simulate a wide set of situations and conditions. The environment has been developed in C, while the visualization module has been developed by means of scripts that use the R environment. The simulations have been performed on personal computers and on clusters of PCs.

The model outlined above is rather complicated and the presently running software adopts some suitable simplifications, in order to test the appropriateness of some mechanisms that have been proposed. The main differences between the model described in section 3 and the current version of the simulation environment are the following:

- updating is synchronous
- the only arithmetic operators are addition and subtraction, without interval constraints
- no vote is given to the relationships with other agents
- production instantaneously adjusts to the total amount of requests (stocks are always adequate)
- at each time step all the agents have a certain probability to create a new artefact (no more than one per agent)
- new recipes can be generated by modifying the input of an existing recipe (without changing the recipe supports) or by applying an evolutionary mechanism to the existing recipes (a genetic algorithm, which allows changes in the recipe supports)
- no new agents are generated.

4.1 Generation of New Goals and New Recipes

In order to introduce new products, agents need to have new goals and to develop the recipes to produce them. Two different methods for devising new goals have been tested: 'imitation' and 'imitation and jump'. We will show in section 5 that only the latter is able to sustain continuous innovation.

In ‘imitation’, an agent is allowed to choose a goal identical to one of the existing products. Actually, the goal is chosen with uniform probability among the known artefacts. In such a way a simple sampling of artefact space is performed: the higher the number of artefacts which lie in a certain region of artefact space, the higher the probability of finding that region. The underlying heuristics is that clusters of artefacts are probably high reward regions. The agents try to build an artefact similar (within a given threshold) to the selected one (which represents the ‘goal’ of the imitation): if they succeed, the new artefact is added to the set of those that are already present.

Imitation alone is unlikely to enhance the agents’ activities, since the new artefacts have no chance to be purchased. Indeed, the name of the artefact created by imitation can be either the same as the target one (and, in this case, the agents that already utilize this name do not have any reason to change their previous suppliers) or different from the name of the target (and, in this case, no agent is able to utilize this new kind of artefact). The only possibility for a new artefact to survive is that another agent creates a new recipe, and chooses the artefact itself to build its own product. This event has a very low probability, and in effect with only imitation the initial agents’ activities drop to zero. A way to avoid this fate is that of allowing the agents to change their suppliers in a random way. If this happens, many agents typically survive (as shown in section 5 below).

Although in this latter case some agents survive, it is observed that the production of new artefacts ends after a finite number of time steps. The problem is indeed related to the definition of new goals: as long as agents try to imitate already existing products, they cannot ‘imagine’ new things. It is therefore appropriate to allow agents to change the way in which the goal is chosen by

- multiplying the name of the artefact, selected by imitation, times a randomly se-

lected value [belonging to a finite interval (min, max)]

- adding or subtracting a random value

In such a way the jump allows the exploration of new regions of artefact space and, as it will be shown in the next section, the combination of imitation and jump allows a continuous generation of new artefacts.

Once a goal has been set, an agent needs to develop the corresponding recipe. Two methods have been devised: exhaustive search and genetics.

In exhaustive search, an agent looks at each of its own recipes and, for each input, it substitutes it with another artefact. The exploration concerns all the artefacts known to the agent, although in the present version not all the combinations of known artefacts are tried: indeed, substitutions in existing recipes are done one at a time. The search halts when the goal has been reached; if this never happens, the closest product is considered and, if its distance from the initial goal is smaller than a threshold, the corresponding recipe is added to the set.

While exhaustive search deals only with the input artefacts, in genetic search the recipe supports are also combined to give rise to new recipes. In particular, the method starts by generating an initial population from a set of recipes (which can either belong to a single agent in the ‘lonely’ mode previously described, or to two collaborating agents). In order to have an initial population, the existing recipes are augmented with some modified versions. These modifications are obtained by changing the input (as in the previous method) with some probability.

The evolution of the population proceeds according to the well-known genetic algorithm. In particular, the fitness function of each individual in the population is a decreasing function of its distance from the goal. The main genetic operators are uniform crossover (which has been modified in order to take into account the existence of recipes of different lengths) and

mutation (note that this term is used here with a meaning which is very different from that used in previous section). Mutation involves here changing one input artefact with the closest one among those known to the agent.

4.2 Initial Conditions

Setting up initial conditions ‘by hand’ is time demanding, so we introduced an automatic ‘initial condition builder’. First of all, raw materials are introduced; afterwards, agents are added one at a time. The recipes of the newly added agent can use the already existing artefacts, producing in this way new artefacts. Networks of this kind are such that older artefacts are likely to be connected to more agents than new one.

The networks generated according to the above-mentioned procedure have the common feature that the unsold artefacts (belonging to the agents added last) cannot provide their owners with any source of activation, and therefore the latter cannot survive, causing a ‘domino effect’ that destroys the entire net. To avoid this collapse we can either provide an external reward region or introduce a modification of the dynamics (e.g. the possibility of changing a provider), which leads to self-sustaining rings, as shown in the next section.

5 Simulation Results

Several simulations have been performed, in order to test the model and to understand some of its behaviours. We do not present statistical results, but rather typical model behaviours. The simulations that are shown below are organized in an order that reflects the introduction of an increasing number of features. In particular, the sequence is the following:

- agents operate without innovation
- agents perform innovations, either in suppliers or in goals and recipes, each agent knowing all the existing artefacts

Fig. 1. Activities vs. time in a net with 10 agents. **a** No self-sustaining loops or reward regions. **b** The stabilizing action of reward regions.

- agents perform innovations, either in suppliers or in goals and recipes, each agent knowing only the artefacts produced by a limited number of other agents

Different innovation methods have been tested. In particular, when innovation in recipes is performed by applying genetic operators to the set of existing recipes, we have considered both the case of 'lonely' agents (which modify only their own recipes) and the case of 'collaborative' agents, which share their recipes in the genetic search of new ones.

5.1 Activation Dynamics without Changes

In this paragraph we show the behaviour of the model when agents do not perform any kind of innovation and never change their providers.

As already stressed, the actual initial condition builder generates networks that are never able to survive without external reward. Figure 1a shows the decrease of activities of a net with 10 agents and no self-sustaining loops or reward regions, whereas figure 1b shows the stabilizing action of a region of reward (the initial situation is the same in both simulations). During

these runs the agents cannot innovate and/or change providers.

5.2 Innovation with Complete Knowledge of Existing Artefacts

In this section we will consider agents which make use of different innovation methods (new goals, new recipes), and which can also modify their suppliers. The distinguishing feature of the different cases described here is that agents have a complete knowledge of all the artefacts produced by the other agents. In most cases we will show the behaviour of lonely agents which operate only on their own set of recipes, but we will also show the results of allowing 'knowledge sharing' between pairs of agents.

5.2.1 Imitation and Change of Suppliers

In this section we will consider the effects of innovations (of different kinds) which agents are allowed to perform. Let us first consider imitation only. As it has already been anticipated, using imitation without change of suppliers eventually leads to the vanishing of agent activities, no matter which method is used to generate new recipes (see fig. 2a). A way to avoid

this fate is that of allowing the agents to change their suppliers in a random way. If this happens, many or all the agents typically survive (as shown in fig. 2b).

Figure 2b also shows that the system is self-maintaining. In absence of reward regions there are two possible explanations for this kind of behaviour: (1) there are self-sustaining loops, and (2) the agents innovate at such a high frequency that they have always sold some artefacts.

In effect, given the particular way of building the initial condition, small changes in the network structure (i.e. the change of providers) can create self-sustaining loops. It is possible to test this hypothesis by deactivating the innovation mechanisms (imitation and change of suppliers) and therefore making the rate of innovations vanish. Figure 3 shows that after such deactivation a situation of stability is achieved, demonstrating the presence of self-sustaining loops. The removal of agent 1 at step 2000 changes the levels of stability but is not able to make the whole system collapse; further analysis could determine whether there is more than one loop or whether agent 1 is located downstream with respect to a loop.

Fig. 2. Agents' activities (no self-sustaining loops or reward regions): with imitation and without change of suppliers (the new artefacts are not purchased, and the agents die) (a), and with change of suppliers the new artefacts have a chance to be purchased, and the network is able to self-maintain (b).

Imitation and change of suppliers show a particular limitation, which is evident in figure 4a: after a while, the number of recipes owned by each agent stops growing. This phenomenon indicates that imitation and change of suppliers do not allow continuous innovation; therefore, imitation is not a true 'novelty generator', and in a few steps all the possible combinations are generated (fig. 4b).

5.2.2 Jump as a Continuous Source of Novelties

Note that the fact that innovations come to a halt, in the previous simulations, should not be attributed to a limitation in the methods used to generate new recipes but rather to the limitations in the available mechanism to devise new goals. Indeed, new recipes can be generated (see section 3) either by modifying the input of an existing recipe or by applying an evolutionary mechanism (e.g. a genetic algorithm) to the recipes owned by the agent. The simulations presented above made use of the former mechanism; however, the results are similar in runs with the second mechanism turned on. The creation of new recipes is not upper limited, and all the numbers $n \in \mathbb{N}^+$ are potentially reachable, but

after initial success we saw that no new artefacts are produced.

Let us then allow agents to 'modify' the chosen goal; so now an agent multiplies, with fixed probability, the name of the artefact, selected by imitation, times a randomly selected value [belonging to a finite interval (0.7, 6)]. In order to appreciate the

influence of this mechanism, it is convenient to compare the behaviour of two simulations, starting from the same initial conditions and differing only in the probability of performing the jump (0 in the reference case and 0.3, respectively). In both simulations the innovation process is stopped at step 3000, in order to check the

Fig. 3. Behaviour of the agents' activities under some external changes; step 0–999: imitation and change of suppliers are active; step 1000–3000: the rate of innovations is set equal to zero; step 2000: agent 1 is removed from the simulation.

Fig. 4. Simulation with imitation and random change of suppliers: number of recipes owned by each agent (a), and total number of artefacts and number of sold artefacts (b); the distance between the two lines is due to the innovations which are not purchased by the agents.

Fig. 5. Comparison between the behaviours with and without the goal jump; on the left the images relative to the simulation without jump, on the right the images relative to the simulation with jump. a, b Diameter of artefacts vs. time. c, d Number of different artefacts' names vs. time.

existence at this time of self-sustaining loops (fig. 5).

The jump introduces very strong effects: the number of recipes owned by each agent during the second simulation is 5 times greater than that of the first simulation, and the number of artefacts present at step 3000 is larger by a factor of 4. Even more interesting, the diameter of artefacts (the distance between the names of the most distant artefacts) and the number of different artefacts' names (the number of different kinds of artefacts present at each step of the simulation) are continuously growing, at least during the simulation time span. The network structure is continuously changing.

Note that the use of the word 'jump' here refers to a change in the value of the variable 'goal' which is not necessarily a small one. In effect, it has been observed (and could theoretically be justified) that the effects of jumps become sensible when their size exceeds the threshold which defines the maximum distance between a given goal and the closest product which can actually be built (i.e., when innovation is tried, a new product which can be produced using the existing recipes and inputs is actually produced only if its distance from the initial goal is smaller than the threshold). On the other hand, a jump coefficient that is too large leads to a basically random exploration of an artefact space that is too vast, which is ineffective.

5.2.3 Response to Perturbations

A powerful way to explore the global properties of a system is that of introducing a slight perturbation and following its subsequent dynamics, which can be compared to that of an unperturbed system. For example, we can consider questions like: What is the system response to perturbations? Is the system robust with respect to changes in the structure of the network? What is the typical system behaviour when one or more of its components (agents, artefacts) disappear? In our case we could:

- remove a recipe from an agent
- remove an entire agent
- remove a group of agents at the same moment

Fig. 6. Comparison between the behaviours with and without the externally enforced elimination of some agents (agents 1, 3, 5, 7, 9 disappear, respectively, at step 50, 200, 500, 1000 and 1300). On the left the images relative to the simulation without eliminations, and on the right the images relative to the simulation with eliminations. **a, b** Diameter of artefacts vs. time. **c, d** Number of different artefacts' names vs. time.

- remove agents at randomly selected moments

Different behaviours to such perturbations are possible: figure 6 shows some examples. Agents 1, 3, 5, 7, 9 disappear, respectively, at step 50, 200, 500, 1000 and 1300. There are several consequences, depending upon the particular agent, the time of fault, the number of agents surviving at the moment of disappearance, the current network structure, and so on. Without describing the details, it is interesting to observe that some perturbations have important consequences (elimination of agent 5 and 9), whereas other perturbations do not. The perturbations affect mainly the number of recipes owned by each agent, the number of different kinds of artefacts and the number of artefacts present at each step of the simulation.

Also the diameter of artefacts is deeply influenced. These changes modify in an irreversible way the system structure (agents, artefacts). However, the global properties described in figure 6 display a strong resilience, i.e. the system recovers after crashes.

5.2.4 Using Genetic Search

Three different strategies have been used in these simulations

- an agent performs genetic search using only its own set of recipes
- an agent tries to reach its goals using only its own recipes; if it fails, it chooses another agent with which genetic search is performed. If the search is successful, both partners own a copy of the successful recipe

- agents always perform genetic search in pairs

When joint search is performed, the chosen partner is the best supplier, i.e. the one which is selling most to the agent.

Some results of the simulation are shown in figure 7. Note that strategies a and b give very similar results, probably because, if each agent knows all the artefacts, the need for joint search is small and the two strategies behave in much the same way. An interesting observation, however, is that if joint search is prescribed, the number of different types of artefacts which are generated is higher than in the previous cases, as is also the number of different recipes per agent. Moreover, using the first two methods one observes a decline in the number of different artefacts,

Fig. 7. Number of different artefact types; **a–c** are described in the text.

Fig. 8. Genetic strategy **b**. Knowledge of only one's suppliers. **a** Number of different artefacts' names vs. time. **b** Fraction of artefacts known to each agent vs. time.

which instead seems to remain fairly constant in the third case.

5.3 Innovation with Limited Knowledge of Existing Artefacts

While complete knowledge of the 'artefacts' is probably useful for an agent, it represents an unrealistic assumption in many real world situations. It is therefore very interesting to explore the behaviour of a system whose agents have incomplete knowledge of the products of other agents. In particular, aspects related to the interaction between agents are expected to be better described in such a framework.

In this case, each agent knows the artefacts produced by some other agents (i.e. its own 'network'). The properties of this network should play a major role in the system behaviour, as we will show in the following sections.

We will consider first the case where an agent knows only its suppliers, and we will then analyze the case where the network of acquaintances is enlarged.

5.3.1 Knowing Only One's Suppliers

We start with the simplest case where an agent knows only the products of its cur-

rent suppliers; the agent tries to innovate by

- defining its goal through imitation and jump (where imitation is performed among the known artefacts)
- finding new recipes by genetics, according to strategy b (i.e. it first tries alone, if it fails it cooperates with its best supplier)

It is interesting to note that, if we had used strategy a (i.e. lonely genetics), the system would have died out, while with strategy b it survives. So in this case the two strategies lead to markedly different results, which is likely to be due mainly to the enlargement of an agent's network which is possible only in the latter case.

Note also that the diversity (i.e. number of different artefact types) grows in time and that the network of each agent grows, without however reaching complete knowledge. This is shown in figure 8b where the fraction of artefacts known to an agent (i.e. the ratio of known artefacts/total artefacts) is plotted versus time. It is particularly interesting to observe that, due to the continuous change which is taking place, even an agent which, at a certain time, knows all the artefacts, is bound to lose this property at subsequent time steps.

5.3.2 Knowing Not Only an Agent's Suppliers

A particularly important aspect is the social network of each agent, and the way in which it changes in time. While in section 2 a sophisticated way to handle this aspect has been described, based upon an evaluation of each potential partner, here we limit ourselves to a simple way to 'go beyond' the knowledge of an agent's suppliers, which has been described above. We will consider the case where an agent knows all the products of its current suppliers and all the products of one or more other agents chosen at random, at each time step, among the set of existing agents (which will be referred to below as 'acquaintances').

In order to disentangle this effect from that of enlarging the set of suppliers, due to the combination of recipes described in the previous section, we limit here genetics to the recipes of a single agent. Therefore, the agent is lonely as far as the generation of new recipes is concerned, but it has a wider social network in that it knows the products of some agents which are not among its suppliers.

We have tested three cases which differ in the number of acquaintances per time step; the system is self-sustaining and the

Fig. 9. Genetic strategy **a** with acquaintances. **a** Fraction of artefacts known to each agent vs. time. **b** Number of recipes owned by each agent vs. time.

number of different artefact types grows and then slowly declines in all cases. The most interesting observation is that each agent soon comes to know all the artefacts, and this situation persists in time, even if at each time step a single new acquaintance is considered (fig. 9). This particular behaviour, which may appear surprising, is due to the fact that, in all the simulations above, there is no hard limit on the number of recipes an agent can simultaneously possess. The only mechanisms for eliminating a recipe are indeed (i) lack of inputs or (ii) lack of sales. What is observed in the previous simulations is that each agent has a large number of recipes, and in such a way the network of suppliers covers the whole set of agents.

In order to further investigate this aspect, we introduced a novel mechanism for eliminating recipes: at each time step, there is a small but finite probability of eliminating the less performing recipe, defined as the one which has the lowest volume of sales in the last 15 steps.

Some very interesting phenomena have been observed in this case. The number of agents which are known to each agent in-

creases with the number of acquaintances per time step N_a , as expected, but even in the case of a very high N_a (e.g. 10 new agents in a population of 30) a distribution of the number of known agents is observed: while there is an agent which knows all the other agents and therefore all the artefacts, there are some other surviving agents which know no more than 40% of the agents.

Similar results are obtained for the number of artefacts that are known to the different agents. In figure 10 relative data are plotted: note that, although in the small N_a case the total number of artefacts and agents declines, the relative number (known artefacts divided by total number of artefacts) actually increases, thus allowing the system long-term survival. Indeed in this case the total number of agents declines, while with high N_a this number remains constant.

It is interesting to consider the number of recipes per agent that, as it may be expected, is a growing function of the number of acquaintances. Even in this case a wide distribution of the number of recipes per agent is observed: in the case $N_a = 1$,

the final (i.e. at $t = 2,000$) number ranges between 4 and 8 recipes, while if $N_a = 15$ this number is comprised between 3 and 15. The analysis of the distribution of recipes (as well as the distribution of known agents and artefacts) is one of the most interesting future directions of research that can be pursued with this model.

Finally, let us consider the number of different artefact types, which is a growing function of N_a (fig. 11). Note that in the graph a maximum number of different artefacts is reached, which is followed by relaxation to a lower number which persists for a long time. This is an unexpected feature, which is found in many different simulations and which needs further investigation.

6 Conclusions

The model under consideration belongs to the class of ABMs, which are well suited to bridge the gap between a careful description of the behaviour of single agents (including their interactions with another or a few other agents) and the system-level features (which result from the interaction of many such agents).

Fig. 10. Genetic strategy **a** with acquaintances: fraction of artefacts known to each agent. **a** One acquaintance at each step. **b** Ten acquaintances at each step.

The relationship with the innovation theory of Lane, Maxfield and co-workers requires that intentionality plays a key role. And indeed in I_2M the agents' intentions are explicitly represented as goals, i.e. products which the agents want to produce and, for this purpose, they try to develop new recipes. This is a major distinguishing feature with respect to many ABMs which rely upon chance as a generator of novelty.

Another distinguishing feature is the role of the agents' social interactions (both of the customer-supplier and of the 'acquaintance' type). Although only preliminary simulations have been performed concerning this aspect, it is apparent from the results in section 5 that the outcomes are deeply affected by the parameters which describe how this social network evolves.

In the present version of the prototype a very simplified way of modifying the social network has been tested, based upon chance, while one of the major future directions of research would also involve intentionality in the change of the social interactions.

It is also interesting to observe that the tests performed so far have shown that the network of customers and suppliers is ac-

Fig. 11. Genetic strategy **a** with acquaintances: number of different artefact types. **a** One acquaintance at each step. **b** Ten acquaintances at each step.

tually able to maintain itself, often forming closed loops, and that it is able to evolve giving rise to 'perpetual' novelties. Therefore, the modelling framework which has been envisaged seems well suited for the exploration of the main goals of innovation modelling in Iscom.

A major line for further development concerns a treatment of the agents' social interactions more closely related to the LM theory. In particular, in a system like ours it will be possible to measure the generative potential of a binary relationship, and to compare the behaviour of agents with different strategies. Some agents will follow the heuristics of favouring relationships with high generative potential (as suggested by Lane and Maxfield), while others will use different heuristics like e.g. 'imitate the most successful agents you know' or 'cooperate with your best customer'. In this way it will be possible to explore under which conditions one heuristic outperforms the others.

The measure of the generative potential may be the product of three normalized quantities, each of which can be determined in our model

- mutual directedness, which can be measured by the vote given to a relationship
- heterogeneity of the agents, which can be related to the differences between their sets of recipes
- aligned directedness, which is measured by the distance between the agents' goals

Moreover, in the future it will be useful to develop suitable tools in order to analyze directly the structure (i.e. characteristic path length, clustering coefficient, topology) of the network which develops, connecting different agents and artefacts.

It will also be interesting to consider possible ways to simplify the model, at least in specific conditions. Indeed, in spite of our efforts to keep things as simple as possible, the intimate relationship with the LM theory led us to a rather complicated

model. Although we do not subscribe to the claim that a model with several parameters can describe almost everything [12, 13], we feel that considering slightly simplified versions of the model might provide useful insights about its behaviour.

Finally, we stress that further improvements of the code are under way in order to make it faster, thus allowing simulations which will involve very many agents, which will allow a better exploration of the emergent properties.

Acknowledgements

We wish to express our gratitude to Andrea Ginzburg, who discussed with us at length the features of this model, since the time of its first conception. We are also indebted to Margherita Russo and to many members of the group working on the Iscom project for many helpful discussions and suggestions. This work was supported by the UE-FET initiative under grant IST-2001-35505 Iscom: The Information Society as a Complex System.

References

- 1 Epstein JM, Axtell R: Growing Artificial Societies: Social Science from the Bottom Up. Cambridge, MIT Press, 1996.
- 2 Gilbert N, Terna P: How to build and use agent-based models in social science. *Mind Soc* 2000; 1: 57–72.
- 3 Cioffi-Revilla C: Invariance and universality in social agent-based simulations. *Proc Natl Acad Sci USA* 2002; 99: 7314–7316.
- 4 Ormerod P, Rosewall B: On the methodology of assessing agent-based evolutionary models in the social sciences. 2nd Brisbane Workshop on Complex Adaptive Systems, 2002; www.paulormerod.com/brisbane3PO.
- 5 Gilbert N: Agent-based social simulation: dealing with complexity, 2005; <http://www.complexityscience.org/NoE/ABSS-dealing%20with%20complexity-1-1.pdf>
- 6 Axelrod R, Tesfatsion L: A guide for newcomers to agent-based modeling in the social sciences; in Kenneth L, Judd KL, Tesfatsion L (eds): *Handbook of Computational Economics, vol 2: Agent-Based Computational Economics*. Amsterdam, North-Holland, in press.
- 7 Maxfield R, Lane D: Foresight, complexity and strategy; in Arthur WB, Durlauf S, Lane D (eds): *Economy as an Evolving Complex System II*. Redwood City, Addison Wesley, 1997, chap 4.
- 8 Lane D, Maxfield R: Ontological uncertainty and innovation. *J Evol Econ* 2005; 15: 3–50.
- 9 Lane D: Artificial worlds and economics (I and II). *J Evol Econ* 1993; 3: 88–107, 177–197.
- 10 Fontana W: Algorithmic chemistry; in Langton CG, Taylor C, Farmer JD, Rasmussen S (eds): *Artificial Life II*. Redwood City, Addison Wesley, 1992.
- 11 Fontana W, Buss LW: The arrival of the fittest: towards a theory of biological organization. *Bull Math Biol* 1994; 56: 1–64.
- 12 Di Gregorio S, Serra R, Villani M: Applying cellular automata to complex environmental problems: The simulation of the bioremediation of contaminated soils. *Theor Comput Sci* 1999; 217: 131–156.
- 13 Serra R, Villani M, Colacci A: Differential equations and cellular automata models of the growth of cell cultures and transformation foci. *Complex Syst* 2001; 13: 347–380.