

# Qualitative Reasoning about Food Webs: Exploring Alternative Representations

**Tim Nuttle (Tim.Nuttle@uni-jena.de)**

Institute of Ecology, University of Jena  
Dornburger Str. 159, 07743 Jena, Germany

**Bert Bredeweg (bert@swi.psy.uva.nl)**

Department of Social Science Informatics, University of Amsterdam  
Roetersstraat 15, 1018 WB Amsterdam, The Netherlands

**Paulo Salles (psalles@unb.br)**

Institute of Biological Sciences, University of Brasilia  
Campus Darcy Ribeiro, 70.343-050 Brasilia DF, Brazil

## Abstract

The trophic interaction is perhaps one of the most fundamental processes in ecology. From such pairwise interactions between organisms and their energy supplies, all ecological communities are built. Here, we explore three alternative qualitative representations of the basic trophic interaction, and discuss their potential to serve as basic building blocks for research and teaching about food chains and food webs.

## Introduction

This paper presents the first installment of our new research agenda to describe basic, qualitative building blocks that can be arranged and rearranged to model a wide variety of ecological situations. Our goal is to build an integrated library of re-usable model fragments that ecologists can assemble to address their own research questions.

The trophic relationship is fundamental to all living beings, and thus a good place to launch our quest. Trophic relationships are commonly represented in food webs of interacting organisms. Food webs are made up of multiple, interconnected food chains, which represent a series of direct linkages of one organism eating another. Despite the relative simplicity of these direct interactions, complex system behavior can emerge from the indirect interactions that result from combining these pairwise interactions in various ways (Begon et al 1996).

Qualitative reasoning can provide interesting insights into the dynamics of food webs because it provides a means of organizing such commonsense, pairwise interactions into a framework that allows system behavior to be analyzed under different scenarios (Bredeweg & Struss 2003). Salles et al. (2003) described a framework for representing pairwise interactions of two populations, but this framework has not been used to model food webs or food chains of more than one trophic interaction. An important goal of the current study was to reorganize their basic pairwise relationship so that any number of interacting populations can be interconnected using the same relationship.

It is important to explore how the basic trophic relationship might best be represented because how interactions are qualitatively represented has important consequences for system behavior. For example, trophic relations can be seen from different perspectives: as numbers of individuals in populations, as transfer and cycles of materials (biomass), and as energy flux. Furthermore, because our goals include laying a groundwork for qualitative reasoning in ecology, it is important to carefully explain the reasons for and potential advantages and disadvantages of modeling choices.

Our objective is thus to explore different qualitative representations of the trophic interactions between two populations. These representations should describe the effects on populations when they are connected by a trophic linkage. Starting from this basic building block, one should be able to represent any food web. Here, we describe three such representations, discuss the advantages and disadvantages of each, briefly describe some essential model output that demonstrates their utility to simulate food webs, and suggest improvements that might be made in the future.

## Models considered

We investigate three model representations that capture various aspects of the trophic relationship. The first considers the effect of trophic interactions directly on numbers of individuals in a given population. The second continues with an ontology focusing on number of individuals in the population, but builds the trophic interaction from basic processes of birth and death that affect populations. Finally, the third uses an ontology focusing on the effect of trophic interactions on the energy intake and metabolism of populations.

We constructed the models using the HOMER qualitative model-building environment (Bessa Machado & Bredeweg 2001.) HOMER is a graphical tool for creating qualitative models that can be simulated by the GARP qualitative reasoning engine (Bredeweg 1992). We inspected simulation results using VisiGARP (Bouwer & Bredeweg

2001). Within the description of each model, we introduce and discuss various modeling concepts and primitives that are used to implement each model.

### Trophic interactions as direct influences on number of individuals

We created this model primarily as a proof of concept to see if the QR engine, GARP, would be able to correctly link together multiple occurrences of the trophic-interaction to simulate a multi-level food chain. Thus, we chose a simple representation of the trophic interaction as the direct effect of one population on another. The classical Lotka-Volterra predation model represents this relationship as:

$$\begin{aligned} dV/dt &= bV - aVP \\ dP/dt &= akVP - dP \end{aligned} \quad (1)$$

where  $V$  is the prey (or “victim”) and  $P$  is the predator,  $b$  is the birth rate of the prey,  $a$  is the encounter rate between prey and predators,  $k$  converts encounters to new predators, and  $d$  is the death rate of the predators (Case 1999). Notice that some causal mechanisms are left implicit: e.g., what causes new prey to be born and predators to die?

**Implementation** We began by defining an entity subtype hierarchy. Properties defined for supertypes are inherited by subtypes. For example, all properties that are defined for *populations* also pertain to subtypes *plants* and *animals*. For all populations, we associated the quantity *number*, representing the number of individuals, to populations. We further associated the quantity space *zlmh* (i.e., *zero*, *low*, *medium*, and *high*) with this quantity to specify the various values this quantity can possess. Such specifications are made in “model fragments” that contain pieces of information that are assembled by the reasoning engine to create a model (employing the compositional modelling approach; Falkenhainer & Forbus 1991). The designation of this particular quantity space was arbitrary – we only needed enough values in the quantity space so that different magnitudes between sizes of populations could be shown. Here, we are not interested in what quantitative landmarks these values might correspond to, although future users might be interested in assigning such meaning to them.

Interactions between entities are defined as applying if they are connected with a specific “configuration”. Thus, we defined the configuration *consumes* to specify the trophic interaction between two populations. The consequences of this configuration when it connects two populations are specified in a model fragment (see Figure 1). The effect of the interaction naturally differs depending on who is being eaten and who does the eating: the predator decreases (exerts a negative influence on, I-) the *number* of prey and the prey increases (exerts a positive influence on, I+) the *number* of the predator. Note that the designation of who is predator and who is prey arises only as a consequence of the direction of the configuration arrow, and is not the result of what particular entities (“animals”, “plants”, “herbivores”, “carnivores”, etc.) are involved. The model fragment also stipulates that it only applies when both the predator and the prey populations have *number* greater than *zero*, i.e., they

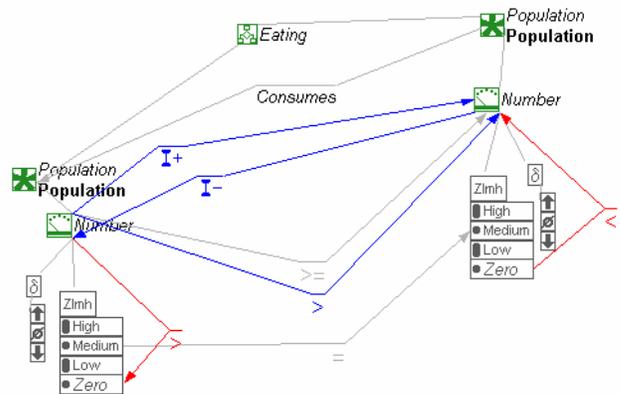


Figure 1: The basic *eating* model fragment specifies the consequences of the configuration *consumes* when it connects two populations. The = between the *medium* points specifies that the two quantity spaces are comparable, and the >= between the two quantity spaces specifies that the consumed population must be ≥ to the consumer population (= is allowed in case the consumed population is *zero*, in which case the consumer must also be *zero*). Depicted is a subtype of this model fragment (indicated by the label *eating* connecting the two populations), which specifies the consequences when both populations are > *zero* (indicated by arrows between value *zero* and its quantity space with inequality signs). When they are both > *zero*, the consumed population is constrained further to be > the consumer population. Influences (I+ and I-) are explained in the text.

“exist”, because a predator obviously cannot consume a nonexistent prey. Similarly, a nonexistent predator cannot consume anything.

Without further constraints, this model would produce a large number of states resulting from the ambiguity of whether opposite influences (I+ and I-) resulted in a net increase or decrease in a population’s size. For example, if a plant is eaten by an herbivore and the herbivore is eaten by a predator, it is ambiguous whether the positive influence of the plant population on the herbivore population would be greater than, less than, or equal to the negative influence of the predator population on the herbivore. So it is unknown whether the herbivore population should increase, decrease, or be stable. In such ambiguous situations, all possible behaviors are produced.

To reduce this ambiguity while at the same time representing the inefficiency of energy transfer (i.e., some energy is lost in its transfer from one trophic level to another; Smith & Smith 1998), we introduce a constraint that stipulates that any population cannot be higher in number than the population it consumes. This included the stipulation that the two quantity spaces are of equal magnitude, though not necessarily of equal value, and thus comparable. We did this by setting the two *medium* point values in the quantity space *number* equivalent for each population (otherwise, *medium* in one population might be equivalent to *low* or *high* in the other). The simulator already knows that *zero* in the two quantity spaces are equivalent. When two point values in a quantity space are

thus fixed as equivalent, this fixes the interval between them (i.e., *low*) to also be equivalent. So the values *zero*, *low*, and *medium* are equivalent for the two populations. We have not included a point above *high* that we could fix, so it is technically impossible to tell whether a population at *high* is greater, equal to, or less than another population at *high*. Because it is impossible to tell the difference, there is no qualitative difference, so the entire quantity spaces are thus made equivalent and thus are comparable (see Figure 1).

Finally, the system to be simulated is specified in a “scenario”. From the scenario, the system determines how it should apply the various model fragments to the various entities and quantities included in the scenario. The initial values are also specified in the scenario. Thus, we defined a four-level food chain with some resource at the base.

**Results and Discussion** The proof of concept that complex scenarios could be simulated by linking together multiple occurrences of the same model fragment was successful; the simulator was able to correctly apply multiple instances of the same model fragment (*eating*) when they occurred together in a scenario depicting a food chain.

Despite its effectiveness as a proof of concept, this model omits some potentially important processes and includes some implicit assumptions that might limit its ability to address some important ecological questions. For example, the only way a population can decrease is if another population consumes it. Therefore, there is nothing limiting the population growth of the highest population in the food chain. Because there is also no other limitation on lower trophic levels, the result is that all populations, if existing, will continue to increase in perpetuity. Additionally, in this model, if a single population existed in isolation, it would neither increase nor decrease in number. Most natural populations have some death rate that is independent of being consumed by another population. It would be beneficial to explicitly represent this nonpredation death process, as is represented by *d* in eq. 1.

Furthermore, it might be more insightful to explicitly represent the other processes affecting each population. For example, the predation process could explicitly represent the effect of both predator and prey population size on predation rate, including how that predation rate affects each population. Similarly, as in the classical Lotka-Volterra model (eq. 1), the growth of populations is currently left implicit. How is it that the prey population causes the predator population to increase? (i.e., birth is not explicitly represented).

Finally, though useful for the present analysis, the constraint on population size not only constrains the values entities can assume but unfortunately also constrains the generality of the model. Hence, it is not possible to investigate phenomena like predator-prey population cycles or response to shifting equilibria. It would be preferable to stipulate these assumptions as *consequences* of specific situations rather than *constraints*. For example, a logical consequence of a predator population being higher than can be supported by its prey would be that the predator population decreases until it reaches an abundance that can be supported.

## Trophic interactions explicitly represented by basic population processes

To make the processes left implicit in the previous representation more explicit, we next created a model representation that considers the consequences of the trophic interaction on basic population processes of birth and death (as discussed in Salles et al. 2003):

$$dN/dt = N(b-d), \quad (2)$$

where *N* is the size of a population in numbers of individuals, *b* is the birth rate, and *d* is the death rate. Here, we consider how these birth and death rates are affected by trophic interactions. Additionally, we included some more ecological realism by specifying some consequences of predator and prey populations being out of balance with each other.

**Implementation** We continue with a similar ontology to that of the previous model, where all entities, including populations and resources, are characterized by the quantity *abundance* with quantity space *zlmh*. *Abundance* is basically the same as number of individuals, but a bit more general.

A consequence of a population having positive abundance is that a *death rate* is introduced (specified in model fragment *death*). The *death rate* decreases (I-) the population *abundance*. The *death rate* can take on two values, *plus* and *zero* (quantity space *zp*), corresponding to the process being active or inactive (“on” or “off”). The *death rate* is positively proportional to the *abundance* because when there are more individuals, more of them die, even if a constant proportion of them die. This is similar to how *d* works in eq. 2.

A model fragment called *consuming* specifies the trophic interaction. This model fragment operates when a *population* is connected to another *entity* with the configuration *consumes* and both the consumer population and consumed entity have positive abundances (Figure 2).

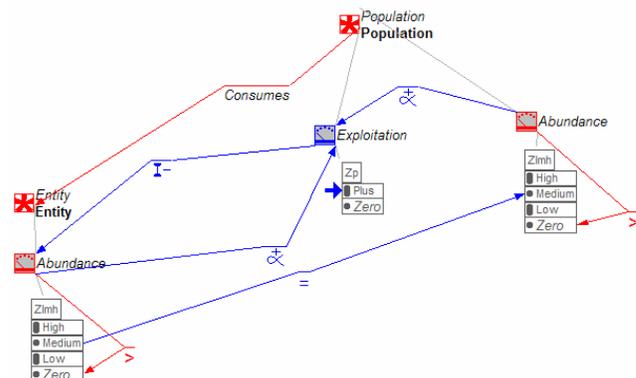


Figure 2: The *consuming* model fragment specifies the consequences of the *consumes* configuration. The quantity *exploitation* is introduced, which is positively proportional to the *abundance* of both the consumer population and the consumed entity. *Exploitation* decreases the *abundance* of the consumed entity; effects of *exploitation* on the consumer are specified in sub-model fragments.

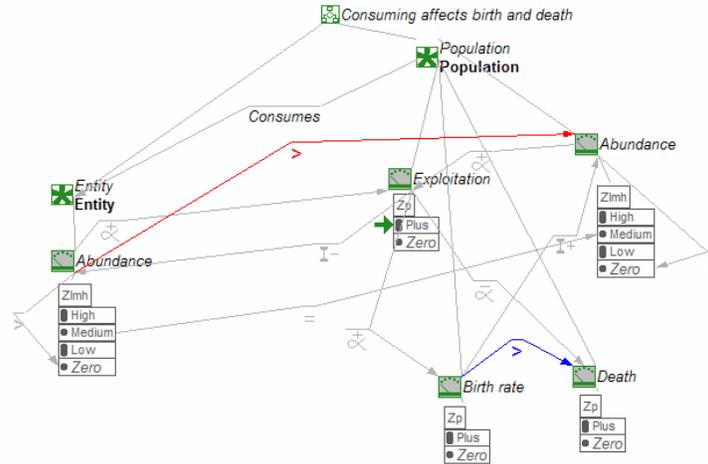


Figure 3: The model fragment *consuming affects birth and death* is a sub-model fragment of *consuming*. Here, the *birth rate* is introduced, which is proportional to the *exploitation rate* and increases (I+) the consumer population’s *abundance*. Also, *death rate* is further specified to be negatively proportional to *exploitation*. A further sub-model fragment, pictured here, specifies that the effect of trophic interaction when the consumed entity is greater in abundance than the consumer population: *birth rate* is than *death*, resulting in a net positive influence on *abundance* of the predator. The consequences of consumed entity’s *abundance* being equal to and less than the consumer’s are specified in similar model fragments.

The two *abundances* are defined as being comparable, so the consequence of their relative sizes can be stipulated (see = sign in Figure 2).

The presence of the *consumes* configuration (Figure 2) results in the introduction of the quantity *exploitation* to the consumer population and the setting of its value to *plus*. We introduce this quantity here because exploitation can only occur in the presence of the trophic interaction. *Exploitation* is positively proportional to *abundance* of both predator and prey (as in eq. 1) and has the effect of decreasing (I-) the *abundance* of the consumed entity.

The effect of the trophic interaction on the consumer population is specified in several sub-model fragments. First, the quantity *birth rate* is introduced, implying that without any food a population will have no births. *Birth rate* is positively proportional to *exploitation*, and has the effect of increasing (I+) the *abundance* of the consuming

population. Additionally, *death rate* is negatively proportional to *exploitation*. Thus there are two counteracting influences on the predator population: *birth* and *death* (as in eq. 2). Further sub-model fragments specify what the net effect of *birth* and *death rate* is. For example, we specify that if the *abundance* of the consumed entity is greater than that of the consumer population, *birth rate* will be greater than *death rate* (Figure 3); hence, there is a net positive influence on population *abundance* of the predator. If the opposite is true, *births* will be less than *deaths* and the net influence will be negative. Finally, if both *abundances* are equal, then *births* will equal *deaths* and the net influence will be zero.

If it is wished to model a system where a resource at the bottom of the food chain is resupplied, rather than used up, the “agent” model fragment *resource control* can be applied, which maintains a constant resource *abundance*. This is accomplished by adding balancing influences (I+ and I-) and setting the derivative of *abundance* to stable; the system calculates the magnitude of these influences to satisfy this constraint. Agent model fragments are only activated when they are specified in a scenario. Thus the abundance of the entity at the base of the food chain can be depleted completely unless *resource control* is specifically invoked.

We specified several scenarios to investigate this representation’s ability to simulate food chains. Due to ambiguity, too many states were generated when more than a few trophic states were included and the agent *resource control* was applied. Therefore, we present here only relatively simple scenarios. First, we investigated a scenario with one population (set to *high*) subsisting on a resource (set to *zero*) to ascertain the basic concept that populations without food should die out. Next, we investigated another

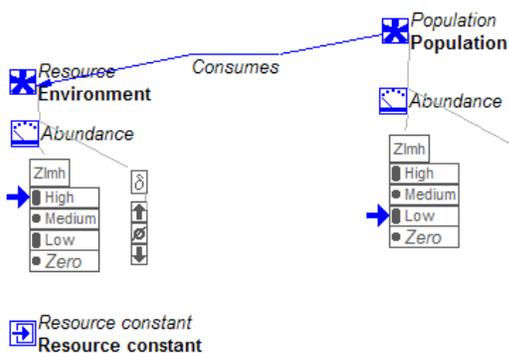


Figure 4: Scenario of one trophic level subsisting on a resource. The agent *Resources constant* was applied in one case and not applied in another.

two scenarios of one population consuming a resource: one where the *resource* was held constant by applying the *Resource constant* agent and one where it was not (Figure 4). Finally, we investigated a scenario where there was an additional trophic level and the *resource* was free to decrease. In each of these cases, *abundance* of the resource was initially set to *high* and of populations to *low*.

**Results and Discussion** The basic test that a population without a resource to consume should die out was successful. When resource *abundance* was held constant at *high* and population *abundance* started out *low*, not surprisingly showed somewhat opposite behavior: with several different pathways possible to this one end state (Figure 5), the population increased in *abundance* to *high*.

When resource *abundance* was not held constant, the simulation was too complicated to adequately display graphically, so we simply describe the most important features. The full simulation produced 43 possible states, including five possible starting and five possible end states. One of these end states was expected: all quantities reached *zero* because the population used up all the resource and then died out. Two of the other end states that were produced appear to coincide with an increasing population and a decreasing resource that should in the next state reach *zero*. However, this creates a “conflict” because if that were to happen, the population *abundance* would want to immediately go from increasing to decreasing without first being stable because of the sudden inactivity of the *consuming* process. This transition is not allowed under qualitative process theory (Forbus 1984), thus creating a conflict and the inability of the simulator to produce a subsequent state. The other two end states have *high* increasing *abundance* for the population and *high* decreasing *abundance* for the resource. This time, the population and resource are not able to reach equal *abundance* and subsequently decline. Thus, there also appears to be a conflict that prevents the expected transitions from occurring.

Although such conflicts reveal that our representation is not perfect, some interesting behavior was produced that allows some ecological insight. For example, the simulation produced an instance where the system cycled around a particular state (population and resource *abundance* both *low*). Eventually, the cycle is broken when all of the resource is used up. However, it is impossible to tell qualitatively how long this behavior would continue. It is interesting to see, however, that this sort of qualitative spiraling or cycling can be produced. Predator-prey population cycles are a well known phenomenon in population biology, and can also be produced under certain conditions by the Lotka-Volterra model (eq. 1; Case 1999).

Further work on this representation would help us better understand why some expected behaviors do not occur, as well as understand why the unexpected behaviors do occur. Additional model refinements making the representation more explicit might then address these problems.

However, we decided not to pursue these refinements with the current ontology. A disadvantage of an ontology focusing on number or abundance of individuals is that this

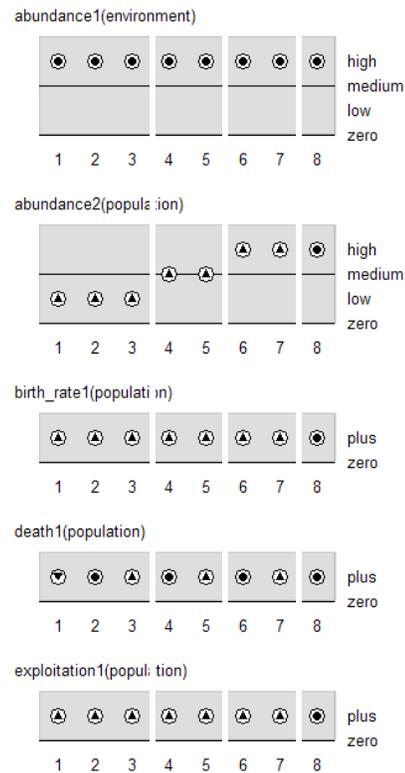
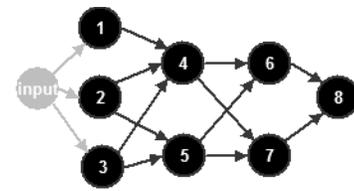


Figure 5: State graph (top) and value history (bottom) for the scenario depicted in Figure 5. There are many possible pathways from the possible beginning states: 1, 2, and 3 can all lead to 4 or 5; 4 and 5 can each lead to 6 or 7; 6 or 7 can lead to 8: a stable population equal in abundance to the resource.

ontology only makes sense when the trophic interaction operates on whole individuals. However, consider herbivory, e.g., a giraffe feeding on acacia trees. The giraffe does not consume entire acacia trees as it browses – an entire population of giraffes would not consume even a single entire acacia tree, and thus the abundance of acacia trees is unaffected by giraffes feeding on them. Nevertheless, giraffes do consume – and in the process influence – acacias.

Furthermore, this ontology is only direct when the sizes of the various individuals are similar. Consider zooplankton consuming phytoplankton. Because zooplankton are considerably larger than phytoplankton, does it make sense to compare their numbers? Moreover, fish that eat zooplankton are orders of magnitude larger than zooplankton, not to mention phytoplankton. Does it make sense to compare numbers of individuals across populations – to say that the influence of some number of phytoplankton

on zooplankton is the same as the influence of the same number of fish on zooplankton? Of course not. A way to finesse this problem is to assume some scaling, where you explain that the comparison points of different populations actually represent different quantitative values (e.g., 1 consumer = 1000 consumed). Thus, the number ontology becomes indirect. Although we are not really concerned with such quantitative landmarks for our present purposes, it is generally desirable to avoid unnecessary confusion where possible. Thus, an ontology based on number or abundance leaves much to be desired in terms of generality.

### Trophic interactions mediated by energy balance in populations

In this representation, we hope to address many of the limitations of the ontology focusing on number of individuals while also clearing up some of modeling conflicts and ambiguities.

Here, we describe the basic trophic interaction between populations using an ontology focusing on energy balance and biomass. The growth of a population's biomass ( $B$ ) depends on the balance between the amount of energy a population assimilates ( $a$ ) and expends from metabolic respiration ( $r$ ), which are each proportional to the biomass (with conversion factor  $k$  of energy to biomass):

$$dB/dt = B(a-r)k. \quad (3)$$

This ontology overcomes the disadvantages of the number ontology. Consider again giraffes eating acacia leaves. Now, it is easy to see that the giraffes remove biomass from the acacia trees, regardless of the number of acacia trees. Conversely, giraffes add biomass to their bodies by eating acacia leaves. If the giraffes do not eat, they will have no assimilation of biomass, while still having metabolic costs that take away their biomass. Thus, the biomass of the giraffe population would decrease because the giraffes starve. Similarly, it is more natural to think in terms of biomass of plankton than the equivalent in terms of millions of individuals. The biomass values are directly comparable, needing no assumed scaling between quantity spaces for populations of different types.

**Implementation** We started with a model fragment that defines a population as having a *biomass*, *energy intake*, and *metabolism*. *Biomass* has quantity space  $zlmh$ . *Energy intake* and *metabolism* have quantity space  $zp$  and are each positively proportional to *biomass*. This model fragment also stipulates that if the population *biomass* is *zero*, so is the *energy intake* and *metabolism*. A sub-model fragment specifies an *existing population* as one whose *biomass* is greater than *zero*.

Another model fragment, called *metabolism*, builds from the *existing population* model fragment. Hence, this model fragment does not apply if the population *biomass* is *zero*. Existing populations have a *metabolism* (*metabolism* set to *plus*). *Metabolism* represents the respiration of living tissue, so the result is that *metabolism* decreases (I-) the population's *biomass* (as in eq. 3).

Because every population has a *metabolism* that decreases

its biomass, the tendency of any population's biomass is to decrease. To offset these biomass losses, a population can assimilate biomass by eating something. Alternatively, an external energy source like the sun may result in assimilation of biomass.

For the trophic interaction, we allow a *population* to subsist on *organic matter*, which is in a superior position to *population* in the entity subtype hierarchy. Another type of *organic matter* (beside *population*) is *dead organic matter*, so we are now allowing the detrital food chain to be represented (Smith & Smith 1998). In a separate model fragment, the concept of *any organic matter available* is introduced, which means any organic matter (alive or dead) that has a *biomass* greater than *zero*. Thus, the model fragment that specifies the trophic interaction consists of an *existing population* connected to *available organic matter* with the configuration *consumes*. The *biomass* of the organic matter and population are defined as being comparable, and the *energy intake* of the consumer population is set to *plus*. Finally, *energy intake* is specified to increase (I+) the *biomass* of the population while at the same time decreasing (I-) the *biomass* of consumed organic matter. These specifications are made in the *consuming* model fragment (Figure 6).

To overcome the previous representation's conflict when the population ran out of food when it was increasing, we introduce a new model fragment. This model fragment implements a boundary condition where if a *population's* food supply is *low* it cannot grow (i.e., derivative of *biomass* must be stable or decreasing). The biological justification for this is that when the food supply is scarce, the population cannot obtain more food than it takes to just maintain the current biomass. The idea is similar to the functional response in population biology (Case 1999). The way this works is the system calculates the influences of *energy intake* to be less than or equal to the *metabolism*. Thus, the population stops growing before the *consuming* model fragment becomes inactive, avoiding the conflict.

To capture the concept that some energy is lost in the trophic interaction, we created a series of model fragments that specify the consequences of a population's *biomass* being greater than, equal to, or less than the *biomass* of what it is consuming. If the population's *biomass* is greater than

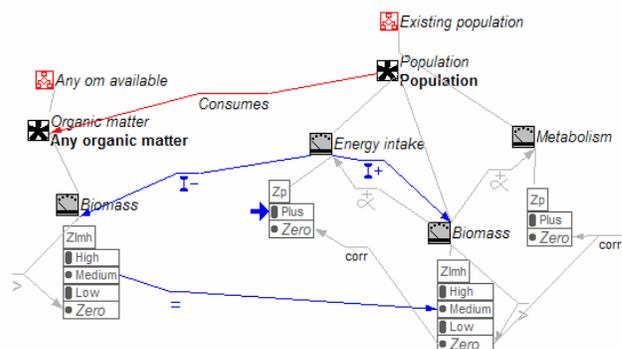


Figure 6: The *consuming* model fragment consists of an *existing population* connected to *any organic matter available* with the configuration *consumes*.

the *biomass* of its food, then the population's *energy intake* is set to be less than its *metabolism*. The converse is true for the less-than condition, and when the two *biomasses* are equal *energy intake* equals *metabolism*. These consequences are specified in sub-model fragments of the *consuming* model fragment. Although similar to the model fragments depicted in Figure 3, here we specify these consequences as "explicit assumptions". An assumption is explicit when it is expressed in a model fragment (or fragments) that contains a specified assumption label. If present in a scenario, the simulator activates those model fragments that contain the label; if the label is not present in a scenario, those model fragments are ignored. Thus it is easy to apply or not apply assumptions when they are explicitly represented in this way. Explicit assumptions are thus activated similarly to agents, except that explicit assumptions generally apply to values of quantities and agents introduce new processes.

The last model fragment specifies how plants obtain energy from the sun. *Photosynthesis* builds from *existing population*, and specifies that the population's *energy intake* will be *plus*, greater than its *metabolism*, and increase (I+) its *biomass* if the attribute *In the sun: Yes* is applied. An "attribute" is a label that can be attached to an entity. When an attribute is applied to an entity in a scenario, the model fragment where the attribute is specified is applied to that entity (and only that entity). Attributes differ from agents because whereas agents apply to any entities in the scenario that are of the same type as or subtypes of those specified in the agent model fragment, attributes only affect the entities they are applied to.

We simulated a series of scenarios all of which were variations of the one depicted in Figure 7. Specifically, we simulated scenarios with two and three trophic levels, with and without the photosynthesis (*In the sun: Yes*) attribute, and with and without the explicit assumption about relative *biomass*. In all of these scenarios, the initial *biomass* of each population was set to *medium*

**Results and Discussion** We start with the simplest scenario: a plant and herbivore without the photosynthesis attribute and without the explicit assumption about relative *biomass* of each population. As expected, because there was no

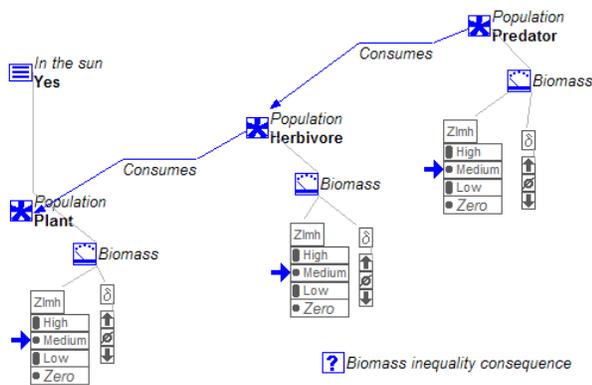


Figure 7: Scenario with a plant, herbivore, and predator. Other scenarios were similar, but with the attribute (label, upper left), explicit assumption (label, lower right), or predator population excluded.

renewable source of energy, only one end state was produced: the *biomass* of both populations became *zero*. However, there were many distinct paths to this one final state. This analysis revealed a problem with this representation: even when the herbivore ran out of plants to eat, its *energy intake* remained *plus* as long as its *biomass* was greater than *zero*. Conceptually this does not make sense and is undesirable. One remedy might be to simply set the *energy intake* to *zero* instead of constraining the derivative of biomass when food is scarce. However, this would negate the function of this model fragment, which is to avoid a sudden (and illegal) switch from increasing *biomass* to decreasing. More thought and experimentation is required to solve this inconsistency adequately.

Next, we simulated the same scenario, but with the photosynthesis attribute applied. From nine initial states and 44 states in the full simulation, there were 15 possible end states, corresponding to different combinations of plant and herbivore *biomass*, including both populations dying out or just the herbivore dying out. Space precludes full exploration and explanation of all pathways and end states.

In the next scenario, we activated the relative biomass

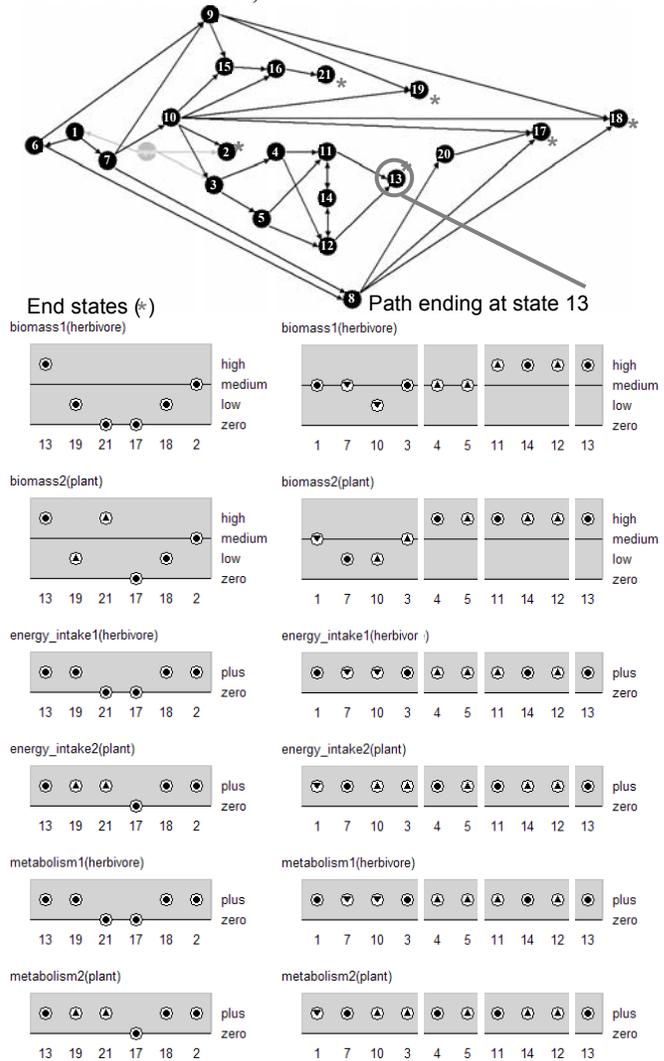


Figure 8: Results of the full simulation of a scenario based on Figure 7, but without the predator population.

assumption for the plant-herbivore system but did not apply the photosynthesis attribute. The full simulations produced seven states and two possible end states; one where both populations were extinct and the other with both populations' biomass at low, the plant decreasing and the herbivore stable. This end state appears to be spurious. Apparently, because in this state the plant and herbivore biomass are equal, the herbivore's energy intake and metabolism are equal (because of the explicit assumption). It seems that the plant should decrease further, removing the equalities and both populations should die out. However, this does not happen, indicating some as-yet non-obvious inconsistency in the relative biomass assumption that needs to be investigated further.

Although the relative biomass assumption appears to be faulty, we continued by applying both this assumption and the photosynthesis attribute to the plant-herbivore system. The results of this simulation are depicted in Figure 8, which highlights the six possible end states (state 19 is similar to the spurious state described in the previous paragraph) and a path to one end state that we found particularly interesting. This path exhibits the biomass spiraling behavior described for the previous representation, but here it occurs in an upward spiral at high biomass (notice two-way arrows at states 11 ↔ 14 ↔ 12).

More complex scenarios involving three trophic levels resulted in far too many states in their full simulations and numerous end states to carefully analyze. However, inspection revealed that they were all basically elaborations of the corresponding two-species system.

## Conclusions

Our three representations of the trophic interaction have varying degrees of success to simulate food chains. We have made much progress in dealing with some of the problems of balancing model realism with managing ambiguity and conflicts, and our representation with the energy/biomass ontology provides a good basis for continued research.

The main challenge is to find insightful ways of dealing with the ambiguity that results from multi-entity systems. In simple systems, like the plant-herbivore system explored in the last section, ambiguity can be insightful. For example, interesting explanations can often be generated for some of the at-first-unexpected pathways and end states the simulation produces. However, the multitude of states and pathways generated by the more complex systems of three species surpassed our ability to easily organize, analyze, and explain.

More work is therefore needed on defining explicit assumptions and attributes that restrict the range of possible behaviors. This reduces generality of any particular scenario investigated but by applying assumptions and attributes only where appropriate, the generality of the framework can be maintained while at the same time making the investigation of particular systems tractable. Thus, with appropriate assumptions and attributes, our approach should allow scaling up from the basic trophic interaction to complex, multi-species food webs. In this regard, the attribute function appears particularly useful: with appropriately

defined attributes, different behaviors can be applied to particular entities in the scenario, allowing fine-tuning of the system to match the characteristics of particular species.

One of the challenges of qualitative reasoning is that such assumptions and attributes need to be built from very simple cause-effect relationships. We have seen that seemingly straightforward constructs, like the relative biomass assumption of the last representation, can produce unexpected conflicts, and even create more complexity than they resolve. We look forward to further refinements that will reduce noninsightful ambiguity and modeling conflicts while also keeping the framework general enough to construct a wide variety of food chains. Particularly, we will continue working on simulating such phenomena as systems that contain multiple predators subsisting on single prey and single prey with multiple predators. These investigations will lead to being able to simulate food webs consisting of multiple interconnected food chains.

## Acknowledgements

We thank the *Deutsche Forschungsgemeinschaft* for support through grant GRK 266 to the University of Jena.

## References

- Begon, M, Harper, JL, Townsend, CR. (1996). *Ecology: Individuals, populations and communities*. Oxford, UK: Blackwell.
- Bessa Machado, V, Bredeweg, B. (2001). *Towards interactive tools for constructing qualitative simulations*. Presented at Proceedings of the 15th International workshop on Qualitative Reasoning, QR'01, San Antonio, Texas, USA
- Bouwer, A, Bredeweg, B. (2001). VisiGarp: Graphical representation of qualitative simulation models. In JD Moore, G Luckhardt Redfield, JL Johnson (Ed.), *Artificial Intelligence in Education: AI-ED in the Wired and Wireless Future*. Amsterdam, The Netherlands: IOS
- Bredeweg, B. (1992). *Expertise in qualitative prediction of behavior*. Ph.D. Thesis. Department of Social Science Informatics, University of Amsterdam, Amsterdam, The Netherlands.
- Bredeweg, B, Struss, P. (2003). Current topics in qualitative reasoning. *AI Magazine* 24: 13-6
- Case, TK. (1999). *An illustrated guide to theoretical ecology*. Oxford, UK: Oxford University Press.
- Falkenhainer, B, Forbus, K. (1991). Compositional modeling: finding the right model for the job. *Artificial Intelligence* 51: 95-143
- Forbus, K. (1984). Qualitative process theory. *Artificial Intelligence* 24: 85-168
- Salles, P, Bredeweg, B, Araujo, S, Neto, W. (2003). Qualitative models of interactions between two populations. *AI Communications* 16: 291-308
- Smith, RL, Smith, TM. (1998). *Elements of Ecology*. San Francisco, California, USA: Benjamin Cummings.