



# **ROMA: A framework to enable open development methodologies in climate change assessment modeling**

Joshua Introne, Robert Laubacher, Thomas W. Malone

*MIT Center for Collective Intelligence*

MIT Center for Collective Intelligence Working Paper No. 2011-03

A revised version of this paper will be published by *IEEE Software*, Nov/Dec 2011

June 2011

**MIT Center for Collective Intelligence**  
**Massachusetts Institute of Technology**  
<http://cci.mit.edu>

**ROMA: A framework to enable open development methodologies in climate change assessment modeling.**

*June 20, 2011*

Joshua Introne, Robert Laubacher, Thomas Malone

Center for Collective Intelligence  
Massachusetts Institute of Technology  
[jintrone,rjl,malone]@mit.edu

**Abstract:** Models play a central role for climate change policy makers, but are often so complex and computationally demanding that experts are required to run them and interpret their results. This reduces stakeholders' ability to explore alternative scenarios, increases perceptions of model complexity and opacity, and can ultimately reduce public confidence in these models.

In this article, we introduce ROMA (Radically Open Modeling Architecture), a web-service that is intended to address these problems. ROMA provides two core functionalities: (1) It supports the creation and running of *surrogate* simulations, which are fast approximations of much larger integrated assessment models, and (2) it offers a componentized view of models and stored model runs and allows clients to combine components to create new executable composite models.

ROMA is currently used to provide modeling functionality in the Climate CoLab. In time, we hope it can support an open development community for climate change models.

**Keywords:** I.6.5.a [Modeling methodologies], I.6.7 [Simulation Support Systems], D.2.11.b [Domain-specific architectures], J.2.e [Earth and atmospheric sciences]

## 1. Introduction

Computational simulation models help support scientifically grounded “what-if” analyses by translating specialized knowledge into tools that can project the likely future impact of current actions. Models have thus become important in a variety of policy domains. In recent years, a number of software platforms for environmental policy-making and urban planning have added simulation models to decision support tools in order to provide stakeholders with direct access to these models, and this trend continues (e.g. [1][2][3]).

Models also play a central role for climate change policy makers, but are so complex and computationally demanding that experts are required to run them and interpret their results, creating a bottleneck between models and stakeholders. This reduces the flexibility that individual stakeholders have to explore alternative scenarios and limits the number of stakeholders that can query the models. It also makes models more opaque to stakeholders because experts summarize model results and omit detail about numerous assumptions made by a model. Complexity and opacity can reduce public confidence in these models.

Drawing inspiration from open source development practices, we seek to address these problems by providing support for modularization of and open access to models that can inform climate policy deliberations. To this end, we have developed a publicly accessible web-service called ROMA (**R**adically **O**pen **M**odeling **A**rchitecture) that allows anyone to create, combine, and run modular simulations. ROMA currently provides the modeling functionality in the Climate CoLab (<http://climatecolab.org>), a collective intelligence application [4] in which large numbers of people work together to develop proposals to address climate change [5]. In time, we hope that ROMA will support a community focused on model development and analysis.

## 2. Design

ROMA was initially developed as part of the Climate CoLab. In the CoLab, community members run models to predict the outcomes of proposals to address climate change. The application of modeling in the CoLab helped crystallize two technical design challenges for ROMA. First, it must simplify complex, computationally expensive models for a broad, web-based community. Models must execute rapidly and without great loss of accuracy, and we must be able to flexibly tune the interface of any model to meet the needs of diverse users.

Our second design challenge was how to use modeling functionality to support collective intelligence. Research has demonstrated that large groups of diverse individuals can find better solutions to problems than similarly sized groups of experts, but only if individuals have a basic understanding of the domain and are free to explore the space [6]. Thus, ROMA needs to provide modeling capabilities that could inject expertise into the users’ exploration of the domain to achieve the necessary level of proficiency, yet still allow individuals to try out ideas not foreseen by model creators.

ROMA provides two core functionalities to meet these design challenges:

1. The service provides tools for generating and running *surrogate* simulations of much larger integrated assessment models (IAMs; integrative models that predict the impacts of climate change across a variety of domains). Surrogate models can be “run” very quickly, and easily customized to reduce the number and complexity of their inputs.
2. The service offers a uniform API and componentized view of models and stored model runs and allows clients to combine components to create executable composite models.

These features make it possible for many stakeholders to explore climate and integrated assessment models directly, and reduce the problems of complexity and opacity by emphasizing modularity and standardizing metadata requirements for models. The design also supports a division of labor in which experts in different specialties can easily add new component models that stakeholders can then combine with others to explore competing assumptions about the world.

## 3. ROMA Service Architecture

Within ROMA, a simulation model is described by its inputs and outputs. Inputs and outputs are called **variables**, can be of any standard data type (e.g. integer, double, text), and can represent vector or scalar values. A model can also be associated with other metadata (such as a name and description) and a URL where the model can be executed. For models that can be run, ROMA publishes its own URL that clients can use to run the model.

A data set that is generated when a model is run is called a **scenario**, and ROMA stores all scenarios that it creates. A scenario contains all of the concrete input and output values from a model run and a reference to the model that generated it. For composite models, the scenario will also contain sub-scenarios that correspond to inputs and outputs of each component model in the composite.

Because it maintains a connection between scenarios and the models that create them, ROMA can semi-automatically update stored scenarios if a model changes. It is also possible to swap out sub-scenarios or replace component models to update a composite scenario. This enables less tightly coupled workflows around the creation of scenarios. For example, a team developing a scenario for a global emissions policy in the CoLab might plug in different national policy scenarios that have been developed separately.

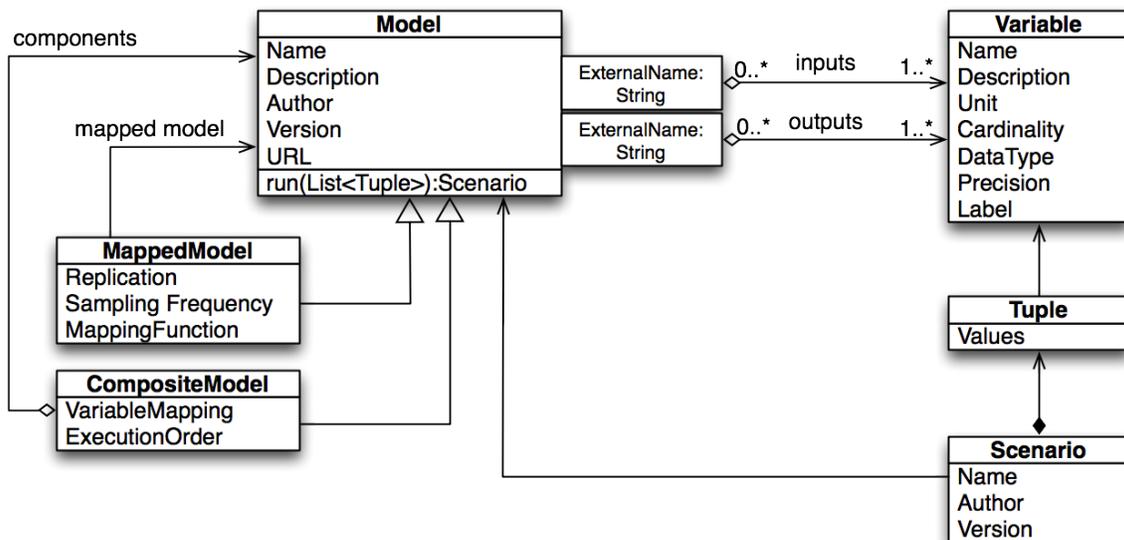
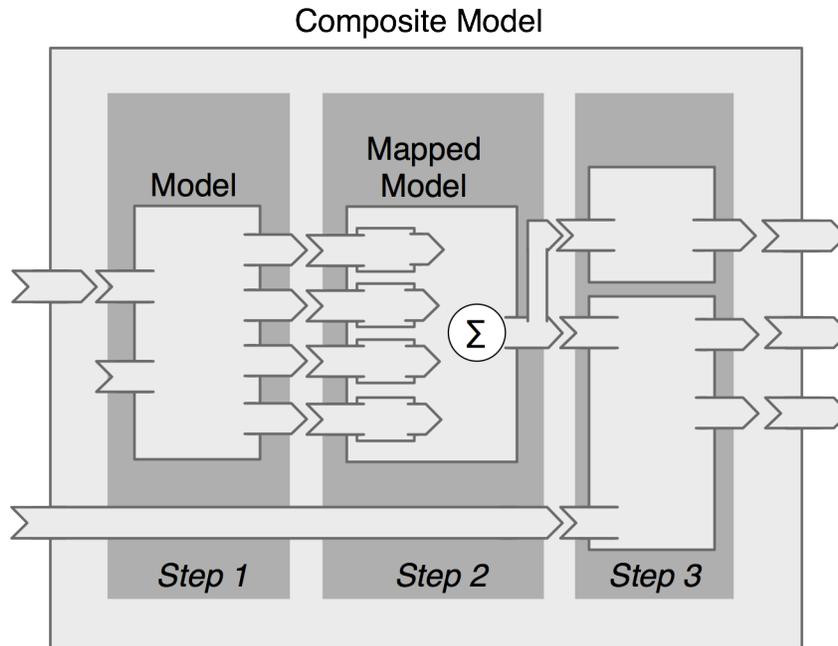


Figure 1: Partial class diagram for the modeling service

The class diagram shown in Figure 1 describes the core components of ROMA. The four main elements – models, variables, scenarios, and tuples – are exposed via a RESTful interface that allows web clients to retrieve XML descriptions of these entities. All other functionality is available through a set of web forms.



**Figure 2: A notional schematic illustrating how models may be connected. A composite model consists of several steps that embed component models.**

ROMA offers two kinds of support for combining models. Mapped models can transform the inputs and outputs of other models, and composite models contain other models and maintain connections between them. Figure 2 illustrates how these types of models can be used in ROMA.

The following transformations are possible using mapped models:

- **Replication** – A model may be repeated  $n$  times over incoming values with higher cardinality. For example, a mapped model with a replication value of  $n > 1$  can be used transform a model that accepts scalar values into a model that accepts vectors with a cardinality  $n$ .
- **Sub-sampling** – The output cardinality of any model may be reduced by sub-sampling its outputs at a given frequency. For example, a model that provides predictions for atmospheric CO<sub>2</sub> for each year over the course of a century may be sampled at a period of 10 years in order to generate data for another model that requires decadal CO<sub>2</sub> values as inputs.
- **Many-to-one mapping** – The output cardinality of any model may also be reduced by applying a many-to-one mapping function (e.g. sum, average, first, last, etc.).

These transformations can be combined, in which case they are applied sequentially as ordered in the preceding list. Thus, a model is first repeated over its inputs, the results of that operation sub-sampled, and finally combined using the many-to-one mapping function if one is specified.

Composite models arrange their component models in a series of ordered, connected steps. Each step can contain any number of models as long as they have no dependencies on each other. A set of connection descriptors specifies connections from the composite model inputs to steps, connections between steps, and from steps to outputs of the composite model. ROMA only allows variables to be connected if have the same units, data type, and cardinality, but more sophisticated compatibility checking is left to the composite model creator. Connections between steps must be from output variables in an earlier step to input variables in a later (though not necessarily adjacent) step so that cycles cannot occur. When a composite model is run, steps are executed in order.

Running a composite model produces a composite scenario that contains references to scenarios generated by each of the component models. It is possible to replace component scenarios (that do not have their inputs determined by upstream models) after a composite scenario has been generated, and the system will

calculate all downstream changes and update the version number of the composite scenario. Similarly, it is possible to replace a component model in a composite model with a new model that has the same inputs and outputs and request that the system update all scenarios to the new composite model.

### *Surrogate Models*

Surrogate models are often used in other domains when the “real” models are too expensive to run for all parameter combinations of interest [7], or model authors prefer to control access to their technology. A surrogate model can be constructed by interpolating between known data points generated by the actual model. In practice, surrogate models are often elaborated as the parameter space of a model is explored. In the case of climate and integrated assessment models though, we generally have access to published data but not the actual models, and so construct surrogate models based on this data.

To simplify the generation of surrogate models in ROMA, the service accepts scenario-based data—a form commonly used for presenting data from IAMs—and automatically generates surrogate models. We are currently developing a user interface that will make it easy for anyone with access to such a set of scenarios (for instance, a model creator) to create a surrogate model within ROMA.

Surrogate models provide users with a very rapid estimate of much larger models for a bounded region of their parameter space. Of course, these estimates are only approximations, the accuracy of which depend on the curve fitting algorithm used, the amount of data available, the complexity of the output surface, and other factors. The tradeoffs between speed and accuracy need to be weighed for each particular application and domain in which surrogate models are used.

### *Model execution and spreadsheet models*

ROMA is intended to work with models that run on other servers. For ROMA to be able to run an external model, it must provide a URL that accepts a form post with values for each input variable in the model and return data to ROMA. ROMA is agnostic with respect to the technology that runs individual models, but no provision is currently made for models that have long execution times (greater than the http request timeout) or require scheduling.

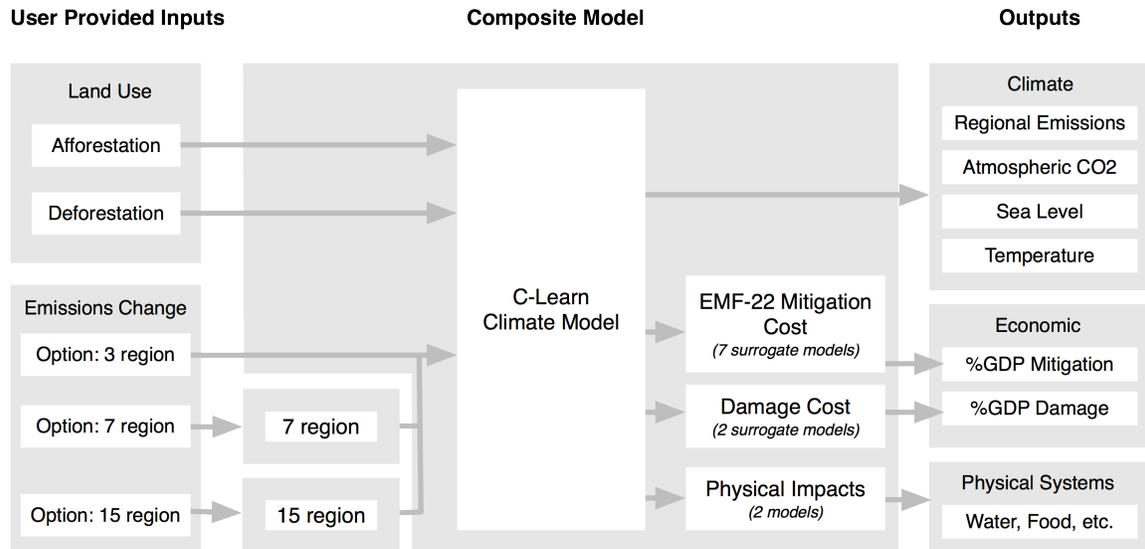
In addition to externally hosted models, ROMA includes functionality that can transform a spreadsheet into an executable model. Spreadsheets are used to implement surrogate models. Spreadsheet models map input and output variables to cells and cell ranges in a spreadsheet. The user defines this association when uploading a spreadsheet to ROMA, and the user-supplied functions that are embedded in the spreadsheet perform all model calculations. The system uses the spreadsheet engine available from the open-source Apache POI project (<http://poi.apache.org/>) to run a spreadsheet model.

Although they are computationally limited, spreadsheets are widely understood and used by many people to create informal models that support decision-making. Thus, spreadsheets provide an easy way to “open up” modeling to a broad community without requiring users to learn a domain specific language for building models.

## **4. Application of ROMA in the Climate CoLab**

In the Climate CoLab, every user proposal must be attached to a ROMA generated scenario that predicts the impacts of that proposal. The CoLab uses the XML data provided by ROMA to generate an interface that allows users to enter input variables, run models, and view stored results.

Thus far, all proposals in the CoLab have been for a global agreement to address climate change, and contributors have used one of three variants of a single composite model to develop scenarios. The composite model combines a climate simulation with models that predict economic and physical systems impacts (Figure 3). The variants of the model differ in the degree of granularity with which users specify emissions reduction commitments for the world’s nations.



**Figure 3: MIT Composite Model inputs, modules, and outputs. The 3, 7, and 15 region inputs for emissions are options in the interface that allow the user to specify emissions reductions at different levels of granularity.**

To run the model, users specify global land use goals and emissions commitments broken out by region as inputs. They may choose to specify emissions targets for three (developed countries, rapidly developing countries, and other developing countries), seven (with larger economies broken out), or fifteen regions. Models in the seven and fifteen region variants of the composite model transform the emissions inputs into the three regions accepted by the C-Learn Climate Model.

Emissions commitments for three regions and land use goals are fed into the C-Learn climate simulation. C-Learn is a version of C-ROADS [8], a lightweight simulation developed by Climate Interactive (<http://climateinteractive.org/>) that can run on personal computers. C-Learn runs as a separate web-service hosted on an internal server. C-Learn produces predictions for a number of indicators including the Climate outputs shown in Figure 3. C-Learn outputs are for each year from 2000 to 2100 inclusive.

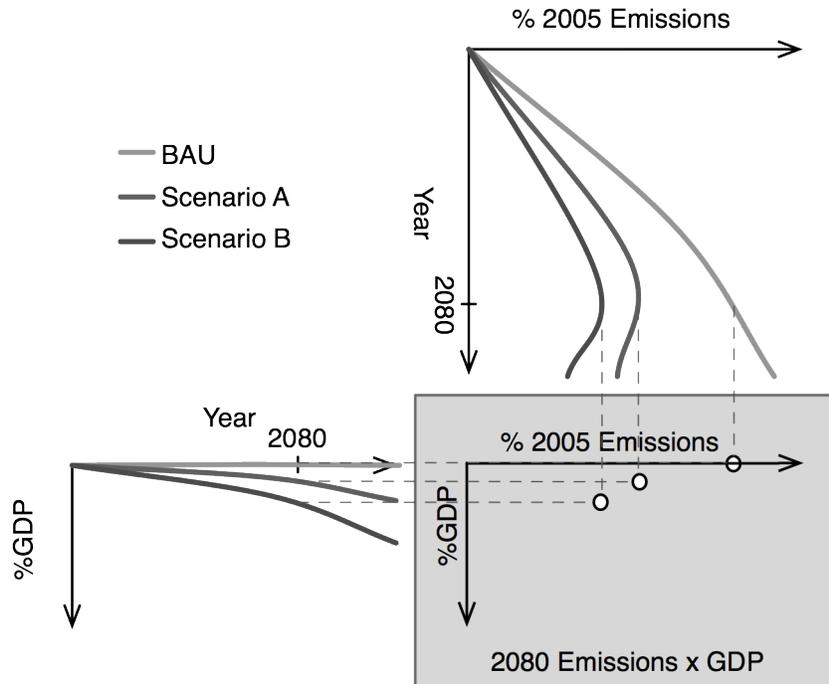
Two physical impacts models produce a brief textual summary of the anticipated effects of temperature change on geophysical (Water, Land, Ecosystems, Singular events) and human systems (Health, Food/Agriculture). Predictions are provided for each Celsius degree of change, taken from [10] and [11]. These models are implemented as a simple spreadsheet model that looks up the appropriate output based on the temperature change by 2100, and a mapped model is used to transform the vector outputs from C-Learn into the scalar output required by the physical impacts models.

Several surrogate models, which are based on data generated by a handful of well-known IAMs, compute economic outputs. These models produce two types of economic costs:

- **Damage cost** – Cost of climate change reported as a percentage deviation from anticipated future GDP if climate change were not to occur.[9], [11]
- **Mitigation cost** – Cost of reducing emissions, also reported as percent deviation GDP from an anticipated baseline. Seven predictions are reported using surrogate models based upon data from the EMF22 exercise [12]. These models are described in more detail in the next section.

#### *Preparing mitigation cost surrogate models*

The mitigation cost models in the CoLab are based on data generated during the Stanford Energy Modeling Forum's EMF 22 exercise. Modeling teams who participated in EMF22 simulated a group of scenarios that reflected a range of potential global mitigation policy approaches plus a reference, or Business as Usual (BAU), scenario with no mitigation policy. Each scenario involved stabilizing greenhouse gas concentrations at a particular target level. Data reported included greenhouse gas (GHG) emissions and sequestration, and a variety of economic indicators including GDP.



**Figure 4:** Development of mitigation cost surrogate; for each year, emissions values in each scenario are plotted against GDP values. The resultant curve is used to infer GDP over the entire range of emissions values for that year.

We created surrogates for the seven models that generated sufficient data. These surrogate models predict the effect of emissions reduction on anticipated GDP for the time span from 2000 to 2100. Change in GHG emissions was chosen as the surrogate model input because emission reductions are the primary mechanism by which GHG stabilization can be achieved and because actions to reduce emissions will be the primary driver of mitigation policy costs.

To construct the surrogates, we used two sets of data for each model: (1) percentage change in fossil fuel CO<sub>2</sub> emissions versus 2005 levels; and (2) percentage reduction in GDP versus the reference (no policy or BAU) scenario. Thus, for each model in each year, we have  $n$  points that associate an emissions level with a percent deviation in GDP, where  $n$  is the number of scenarios used in our analysis for that model (see Figure 4).

To determine the impact on GDP for any emissions level in a particular year, we locate that point on a curve that best fits the  $n$  data points available for that year. We used linear piecewise interpolation to approximate this curve; more sophisticated approaches (e.g. higher order polynomials or spline interpolation) are possible but we did not feel they were justified without additional data. If emissions levels are lower than the most aggressive scenario in a particular year, the surrogate model does not report a value. If emission levels are higher than BAU (the scenario for which mitigation cost is zero), the model simply reports zero percent change in anticipated GDP.

For policy proposals in the CoLab that are too aggressive to be modeled by a particular surrogate mitigation model (e.g. emissions levels are too low in a particular year), the system reports that the modeling team in question likely judges the policy scenario to be technically or economically infeasible.

Some inaccuracies arise for the mitigation cost models in the CoLab because emissions values generated by C-Learn and used as inputs to the surrogate models do not correspond in every detail with the emissions values used in the original mitigation models. For example, CoLab users can specify land use goals to manipulate emissions levels, but the surrogate models do not incorporate this as a source of emissions. Land use accounted for approximately 8 percent of total CO<sub>2</sub> emissions in 2010, so differences in land use

policy can have an incremental impact on both environmental and economic outcomes. To enhance the accuracy of the system, we are exploring the incorporation of land use emissions in the surrogate model in the future.

Despite these limitations, we feel that the level of accuracy offered by the mitigation cost models is sufficient to facilitate the collective problem solving activities and dialogue in the Climate CoLab.

## **5. Future Work**

### *Enhancing the Climate CoLab*

The modeling functionality offered to Climate CoLab users is only a subset of what is possible with ROMA. We plan to introduce more advanced functionality as we develop organizational processes to help scaffold its use. Over the course of this year (2011), the CoLab expects to launch a series of contests to create both national and global proposals for emissions reduction. National and global contests will occur in parallel, and will be phased with interim evaluations of proposals at the end of each phase. We expect these contests to use ROMA's scenario recombination functionality to allow authors of global proposals to select and connect scenarios generated in support of national proposals.

Before launching national level contests, we expect to solicit the development of national models to predict how different energy policies and technology choices will impact emissions and GDP at the national level. To facilitate this process, we hope to expose the surrogate and spreadsheet model creation functionality to model creators via web interfaces. After we gain experience with our users and iterate the design, we expect to make this interface available to a broader community and invite broad participation in the creation of models.

### *Enhancing ROMA*

In our vision of radically open climate and integrated assessment modeling, a large community of modelers and stakeholders will work to develop modules that can be combined to create larger, more sophisticated models, much like an open source software development community. As with open source software, it will be important to develop a community process and tools to help:

- 1) Control for the quality and validity of individual models.
- 2) Determine if modules can be combined.
- 3) Determine if one module can replace another.

Unlike with software, individual models do not have unit tests to help guarantee their validity or integration tests to evaluate their compatibility. A technically perfect model might rest on completely ungrounded assumptions, and any combination of models might be inconsistent if assumptions conflict.

Within the CoLab, the validity of the models has been established via a centrally administered review process with a panel of experts. To support the vision of an open-modeling community, we hope to design processes and technical support to better leverage the collective intelligence of the community. For instance, the community could be invited to look for obvious errors (infinite or impossible values at the extrema of the input space) for each model. Model creators and experts might attach key assumptions to individual models, and experts could weigh in on the validity of those assumptions. These assessments could be summarized to provide policy makers with indications to about model maturity and uncertainty.

The creation and adoption of a set of standard indicators for integrated assessment modeling (e.g. <http://www.epa.gov/climatechange/indicators.html>) will further help support model composition and comparison. Such standards have not yet been created, but we believe that this will occur in the future and hope that approaches like ours could encourage this kind of standardization.

The hurdles to creating community processes around model creation, analysis, and validation are as much social and organizational as they are technical. Integrated assessment models have traditionally been

implemented as monolithic software projects and developed by small teams of experts, and these development processes have led to the complexity and opacity that currently cause difficulties. By emphasizing modularity and offering a set of features that allow stakeholders to become more directly involved in climate and assessment modeling, we hope ROMA can enable social and organizational processes that ultimately improve our chances of creating solutions to climate change.

## References

- [1] M. Matthies, C. Giupponi, and B. Ostendorf, "Environmental decision support systems: Current issues, methods and tools," *Environmental Modelling & Software*, vol. 22, no. 2, pp. 123-127, Feb. 2007.
- [2] I. S. Mayer, E. Van Bueren, P. Bots, H. van der Voort, and R. Seijdel, "Collaborative decisionmaking for sustainable urban renewal projects: a simulation-gaming approach," *Environment and Planning B: planning and design*, vol. 32, no. 3, p. 403-423, 2005.
- [3] B. Friedman et al., "Laying the foundations for public participation and value advocacy: interaction design for a large scale urban simulation," in *Proceedings of the 2008 international conference on Digital government research*, Montreal, Canada, 2008, pp. 305-314.
- [4] T. W. Malone, R. Laubacher, and C. Dellarocas, "The Collective Intelligence Genome," *Sloan Management Review*, vol. 51, no. 3, pp. 21-31, Spring.
- [5] J. Introne, R. Laubacher, G. Olson, and T. Malone, "The Climate CoLab: Large scale model-based collaborative planning," in *Proceedings of the 2011 Conference on Collaboration Technologies and Systems*, Philadelphia, PA, 2011.
- [6] L. Hong and S. E. Page, "Groups of diverse problem solvers can outperform groups of high-ability problem solvers," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 101, no. 46, p. 16385, 2004.
- [7] D. Gorissen, I. Couckuyt, P. Demeester, T. Dhaene, and K. Crombecq, "A surrogate modeling and adaptive sampling toolbox for computer based design," *The Journal of Machine Learning Research*, vol. 11, p. 2051-2055, 2010.
- [8] T. Fiddaman, L. Siegel, E. Sawin, and A. Jones, John, *C-ROADS Simulator Reference Guide*. 2011.
- [9] P. W. D. Nordhaus, *A Question of Balance: Weighing the Options on Global Warming Policies*. Yale University Press, 2008.
- [10] M. L. Parry, O. F. Canziani, and J. P. Palutikof, "Technical summary. Climate Change 2007: Impacts, Adaptation and Vulnerability. Contribution of Working Group II to the Fourth assessment Report of the Intergovernmental Panel on Climate Change.,," in *Report of the Intergovernmental Panel on Climate Change*, M. L. Parry, O. F. Canziani, J. P. Palutikof, P. J. van der Linden, and C. E. Hanson, Eds. Cambridge, UK: Cambridge University Press, 2007, pp. 23-78.
- [11] N. H. Stern, *The economics of climate change: the Stern review*. Cambridge, UK: Cambridge University Press, 2007.
- [12] L. Clarke, J. Edmonds, V. Krey, R. Richels, S. Rose, and M. Tavoni, "International climate policy architectures: Overview of the EMF 22 International Scenarios," *Energy Economics*, vol. 31, no. 2, p. S64-S81, Dec. 2009.