GenePath

http://genepath.org

Use GenePath to:

- Enter, update and maintain your experimental data.
- Manage prior knowledge on genetic relations.
- Infer relations between genes and outcomes.
- Construct genetic networks from experimental data and prior knowledge.
- Obtain explanations about relations in genetic networks.
- Add notes and graphical material to your projects.
- Manage and save genetic data analysis projects in XML.

An Intelligent Assistant to Genetic Analysis

Genetic networks are often used in the analysis of biological phenomena. In classical genetics, they are constructed manually from experimental data on mutants. To provide assistance in construction of genetic networks from experimental data and support in genetic data analysis we have designed GenePath.

GenePath is an intelligent assistant that automates the analysis of genetic data. GenePath employs expert-defined patterns to uncover gene relations from the data, and uses these relations as constraints in the search for a plausible genetic network. GenePath formalizes genetic data analysis, facilitates the consideration of all the available data in a consistent manner, and the examination of the large number of possible consequences of planned experiments.

To make the results of the analysis communicable, and assist in the understanding of underlying principles, GenePath also provides an explanation mechanism that traces every finding to the pertinent data.

In addition, GenePath may propose additional experiments that could refine the network and relate yet unrelated genes.

Implemented as a web-based assistant, and provided with a set of example data sets and networks, our desire was to make GenePath easy to use yet powerful to address real and complex problems in biology and genetics.

GenePath was inspired by the data analysis problems that emerged from studies of Dictyostelium Discoideum, a soil amoeba with a social behavior. The picture above shows spiral formation of Dictyostelium during aggregation.

About this Guide

This is a guide for all those that would like to quickly get an idea on what GenePath is about and how it can be used. It is not a reference guide, as GenePath is easy enough to use without detailing all of its menus and options. Instead, we rather provide a walk-through using a experimental data set on aggregation of *Dictyostelium discoideum*, a soil amoeba with an interesting social behavior.

Aggregation of *Dictyostelium* is just one of many projects that are freely accessible on GenePath's web pages, and you can upload and run them on our GenePath server at http://genepath.org. Browsing though this guide and with using GenePath on-line may be the best way to start learning about it and using it fruitfully.

Inside:

GenePath

Running GenePath

GenePath is easy to run. Just use your favorite browser, visit GenePath's homepage at http://genepath.org and click on a link that says "Run GenePath". This will open a new window — make sure your browser does not block this — where you can start a new project, upload a project from your computer or play with some of the projects that are available on GenePath's server.

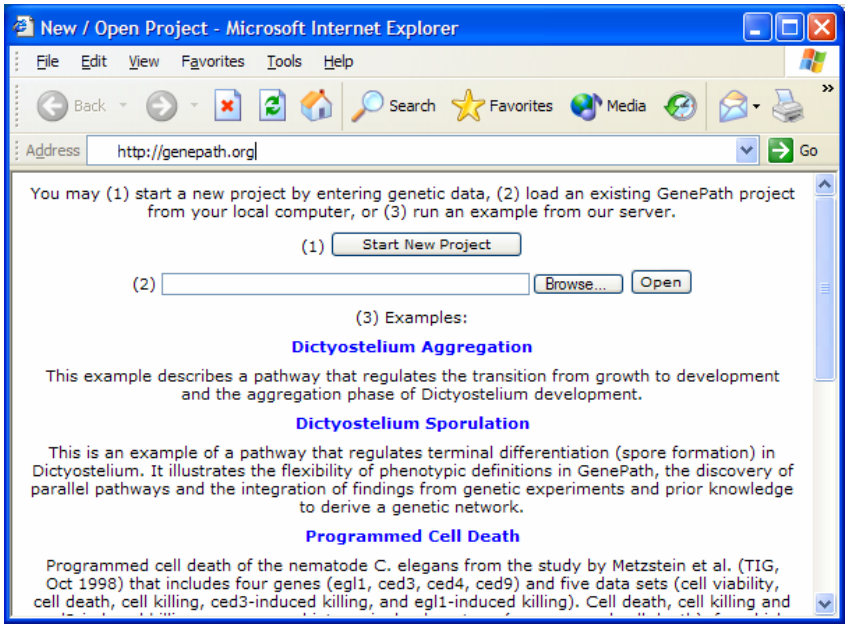

Let us start here by opening a project on aggregation of *Dictyostelium discoideum*. To do this click on the line with "Dictyostelium Aggregation" from the list of examples. The content of the window remains almost the same except for the upper part where a menu bar with three rows of buttons appears. This is the main menu of GenePath with buttons that provide the various functions supported in the program.

The Main Menu

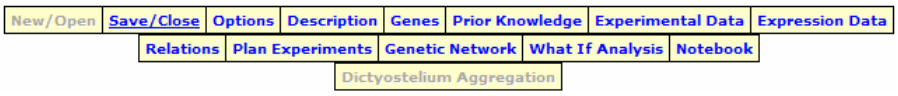

The main menu of GenePath consists of three rows of buttons.

Buttons in the first row are for managing your projects — you can start a new or open an existing project and save your data to a local file. The third button is for setting various options that GenePath's inference algorithms and computational routines depend on. Other buttons in this row will show the forms through which you will define your project and where you will list the genes, prior knowledge, experimental and expression data.

The second row of buttons is about data analysis. GenePath can show a list of relations between genes and outcome, show a resulting genetic network, provide an interface to what-if analysis and help to plan experiments. Notebook is where you can store comments on the project, and is designed so that you can include figures or any graphical material of interest. If you have pictures of the organism you study, or of mutants that are listed in experiments, this is the place to include them.

The last row of buttons shows the names of the projects that are open. Since we have just started with GenePath and opened a single project, there is a sole button on this row. Try opening another project (click on another example project in the list provided) and see how this line changes.

Starting With GenePath

Project Description, Genes, Biological Process, Phenotypes

In GenePath, projects are defined through the list of experiments and prior knowledge. But before these can be entered, you need to give the project a name, define the genes that will be used in the experiments, name the biological process studied in the experiments and define the corresponding list of phenotypes. All this is done on the project description form, and one gets there by pressing the button "Description" from the main menu. A snapshot on the right shows how these items where defined for the project "Aggregation of *Dictyostelium*".

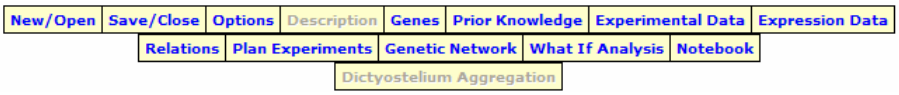

Few comments on the phenotype are in order. First, notice that the phenotypes are qualitative. They would most often be stated using some descriptive names, or by using signs like in the case for our example project. Names of the phenotypes are descriptive and do not have any semantics to GenePath, but the order in which the phenotypes are listed does. The qualitative values should be listed from lowest to the highest value, with the value of the wild-type phenotype (*i.e.*, the phenotype of an organism with no mutations) usually somewhere in the middle of the list.

When using GenePath, bear in mind that this is a web-based program: make sure you press the update button when switching from the project description screen to some other screen. Also, save the projects frequently to prevent the loss of data.

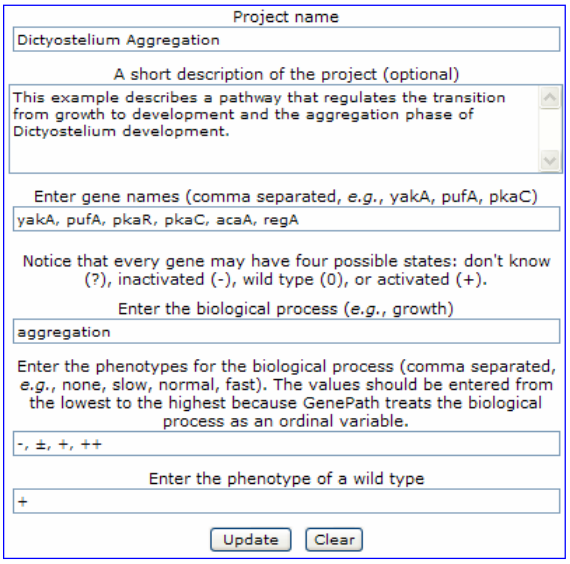

"When using GenePath, bear in mind that this is a web-based program: make sure you press the update button when switching from the project description screen to some other screen. Also, save the projects frequently to prevent the loss of data."

List of Genes

The list of genes can also be accessed through the "Genes" button. Here, you may also provide a comment for each gene in the project. Notice that to edit the information on the particular gene, you need to press on the letter "E" in the row for that gene. Similarly, press the letter "D" to remove a particular gene from the list.

Be careful when you remove a gene which you have already defined and used in your experimental data: removing a particular gene will remove all experiments in which this gene was involved.

To add a new gene, press on the line above the table with a list of genes.

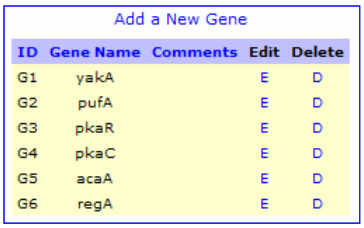

GenePath

Experimental Data

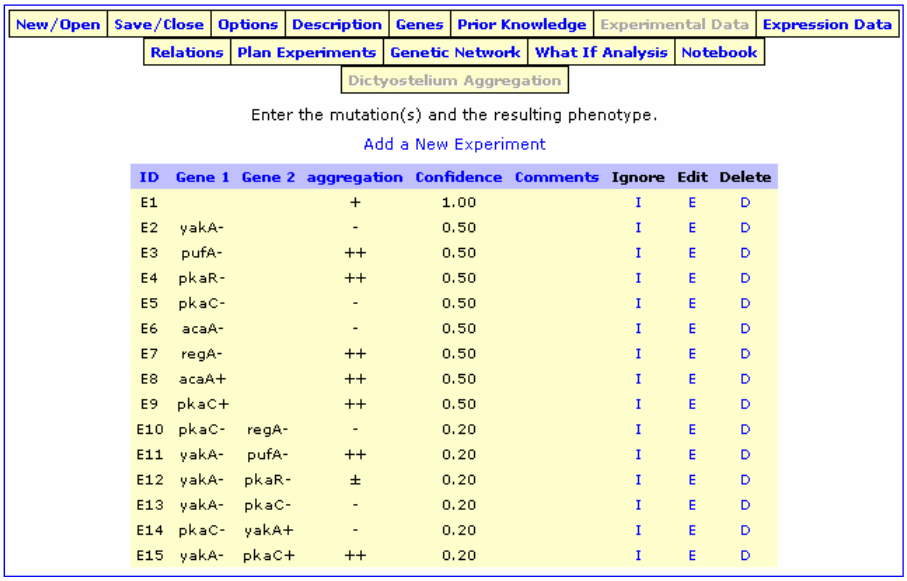

ID	Gene 1	Gene 2	aggregation	Confidence	Comments	Ignore	Edit	Delete
E1			+	1.00		I	E	D
E2	yakA-		-	0.50		I	E	D
E3	pufA-		++	0.50		I	E	D
E4	pkaR-		++	0.50		I	E	D
E5	pkaC-		-	0.50		I	E	D
E6	acaA-		-	0.50		I	E	D
E7	regA-		++	0.50		I	E	D
E8	acaA+		++	0.50		I	E	D
E9	pkaC+		++	0.50		I	E	D
E10	pkaC-	regA-	-	0.20		I	E	D
E11	yakA-	pufA-	++	0.20		I	E	D
E12	yakA-	pkaR-	±	0.20		I	E	D
E13	yakA-	pkaC-	-	0.20		I	E	D
E14	pkaC-	yakA+	-	0.20		I	E	D
E15	yakA-	pkaC+	++	0.20		I	E	D

GenePath derives the relations between genes and outcome from the experimental data that is provided in the project. The snapshot above shows the first few experiments defined in our running example project.

First, notice that the first experiment (E1) does not include any mutation — this is a wild type experiment. Experiments from E2 to E9 are single mutants. Knockouts are marked with minus signs (*e.g.*, yakA-) and overexpressions with plus signs (*e.g.* pkaC+). The last six experiments in the list are double mutants.

You may assign confidence to each experiment in the database. GenePath provides for default confidences that are assigned based on the type of experiment and mutations involved. Alternatively, you may override this by entering some particular value for confidence of your choice.

Experiments may be ignored by pressing "I" in the corresponding row. This will keep the experiment in the data base but ignore it in the consequent data analysis and inference of genetic network. Ignored experiments will have "I" printed in red.

Adding a new experiment or editing the existing one will bring a dialog on the top of the experiment table. For instance, the snapshot below shows this dialog when editing the second experiment in our list. Notice that the experiment may be changed completely (new genes, mutation types, etc.), and you may decide to either update the old experiment accordingly, or add a new experiment to the list. To enter the confidence of your choice, instead of using the default value, unclick the corresponding check box.

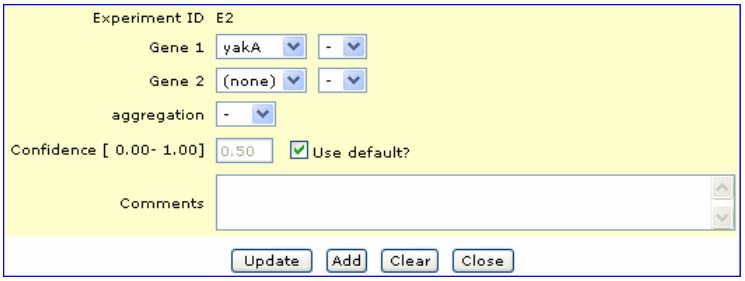

Starting With GenePath

Prior Knowledge

Prior or background knowledge in GenePath is knowledge on any relations between genes or genes and biological process that are known or hypothesized from sources other than your experimental data. For instance, if the particular relation between two genes has already been studied and reported in the literature, and you may want to simply add it as a prior knowledge to your project.

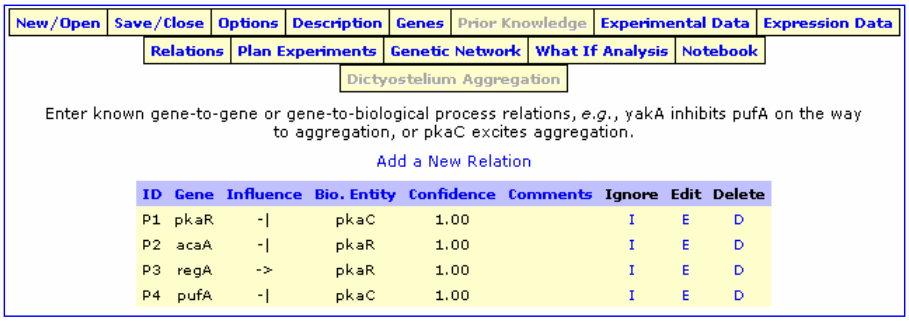

ID	Gene	Influence	Bio. Entity	Confidence	Comments	Ignore	Edit	Delete
P1	pkaR	-|	pkaC	1.00		I	E	D
P2	acaA	-|	pkaR	1.00		I	E	D
P3	regA	->	pkaR	1.00		I	E	D
P4	pufA	-|	pkaC	1.00		I	E	D

The relation between genes and processes in GenePath are either an excitation or inhibition. GenePath marks them with either "->" or "-|", respectively. The relations you may add in the list of the background knowledge are then either on excitation or inhibition between two genes or a gene and a biological process.

The snapshot above shows a list of four relations that are included in the project on the aggregation of *Dictyostelium*.

The snapshot below shows the edit dialog that opens when editing the third relation in the background knowledge list. Just like any edit dialog in GenePath, you may change the data, assign different confidence to the data you are providing, and then either update the edited entry or add a new relation to the list.

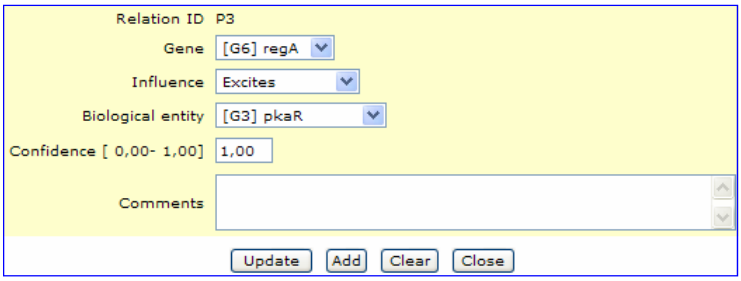

You may consider adding comments to the relations you list in the background knowledge — these, for instance, may point to the particular references where the relation was reported.

Expression Data

Another kind of data that GenePath can use when constructing the network is expression data. Expression data reports on the expression of a particular gene under a specific background, *i.e.*, observed with some mutant. For instance, on the right is an expression data from the Dictyostelium Communication Adhesion project, where, for example, the first row reports on the expression of comC in the lagC– mutant. In the expression data, a gene may be reported as overexpressed (+), underexpressed (-) or expressed normally as in the wild type (0).

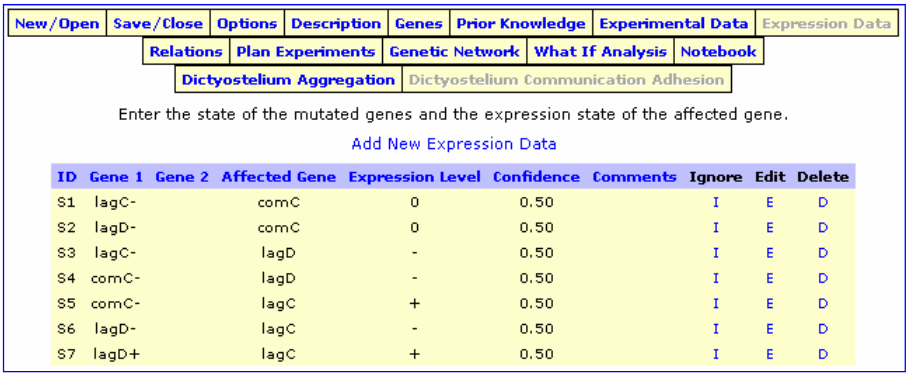

ID	Gene 1	Gene 2	Affected Gene	Expression Level	Confidence	Comments	Ignore	Edit	Delete
S1	lagC-		comC	0	0.50		I	E	D
S2	lagD-		comC	0	0.50		I	E	D
S3	lagC-		lagD	-	0.50		I	E	D
S4	comC-		lagD	-	0.50		I	E	D
S5	comC-		lagC	+	0.50		I	E	D
S6	lagD-		lagC	-	0.50		I	E	D
S7	lagD+		lagC	+	0.50		I	E	D

GenePath

Inference on Relations and their Explanation

"Explanation in GenePath means tracing any finding back to the original set of experiments and reporting how they were combined to make a particular hypothesis on relations between genes."

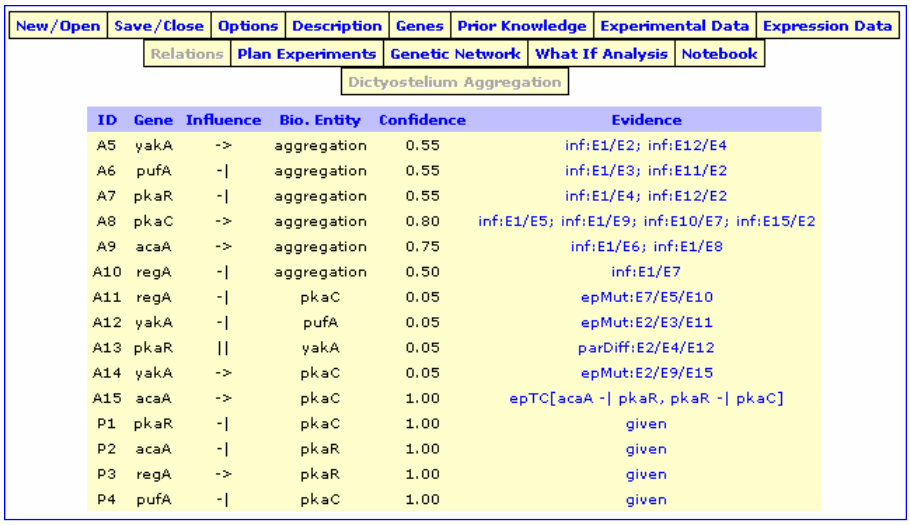

ID	Gene	Influence	Bio. Entity	Confidence	Evidence
A5	yakA	->	aggregation	0.55	inf:E1/E2; inf:E12/E4
A6	pufA	-|	aggregation	0.55	inf:E1/E3; inf:E11/E2
A7	pkaR	-|	aggregation	0.55	inf:E1/E4; inf:E12/E2
A8	pkaC	->	aggregation	0.80	inf:E1/E5; inf:E1/E9; inf:E10/E7; inf:E15/E2
A9	acaA	->	aggregation	0.75	inf:E1/E6; inf:E1/E8
A10	regA	-|	aggregation	0.50	inf:E1/E7
A11	regA	-|	pkaC	0.05	epMut:E7/E5/E10
A12	yakA	-|	pufA	0.05	epMut:E2/E3/E11
A13	pkaR	||	yakA	0.05	parDiff:E2/E4/E12
A14	yakA	->	pkaC	0.05	epMut:E2/E9/E15
A15	acaA	->	pkaC	1.00	epTC[acaA -| pkaR, pkaR -| pkaC]
P1	pkaR	-|	pkaC	1.00	given
P2	acaA	-|	pkaR	1.00	given
P3	regA	->	pkaR	1.00	given
P4	pufA	-|	pkaC	1.00	given

Now for the most interesting part. GenePath uses the data you enter for the analysis. It applies a set of inference patters to the data, and can report results as relations between genes and outcomes. For those methodologically inspired: the particular procedure GenePath uses for logical inference is called abduction, and various techniques from the field of artificial intelligence are used to reason with patterns, relations, and organize them in the genetic network.

The set of relations that GenePath inferred from the aggregation data are given in the snapshot below. Notice that every relation includes an evidence, that is a list of proofs each including a set of experiments that were used to infer the relation. Abbreviations in evidences like "inf" and "epMut" are short names of the patterns that were used.

Clicking on an evidence opens a window with a detailed explanation how a particular relation was inferred. For instance, the snapshot below gives an explanation of the relation "yakA excites pkaC" that was inferred by GenePath. Explanation in GenePath means tracing any finding back to the original set of experiments and reporting how they were combined to make a particular hypothesis on relations between genes.

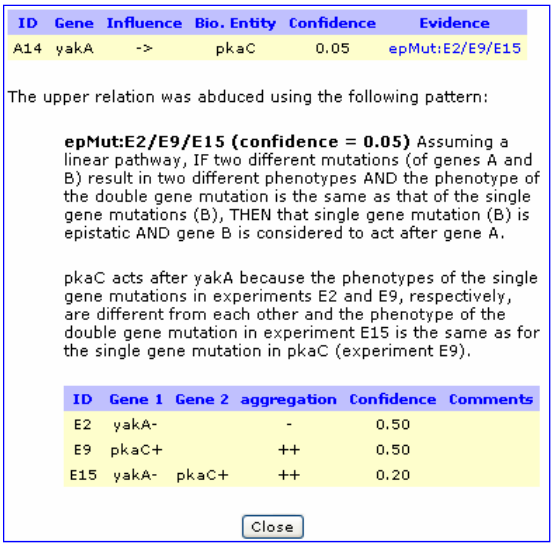

ID	Gene	Influence	Bio. Entity	Confidence	Evidence
A14	yakA	->	pkaC	0.05	epMut:E2/E9/E15

The upper relation was abduced using the following pattern:

epMut:E2/E9/E15 (confidence = 0.05) Assuming a linear pathway, IF two different mutations (of genes A and B) result in two different phenotypes AND the phenotype of the double gene mutation is the same as that of the single gene mutations (B), THEN that single gene mutation (B) is epistatic AND gene B is considered to act after gene A.

pkaC acts after yakA because the phenotypes of the single gene mutations in experiments E2 and E9, respectively, are different from each other and the phenotype of the double gene mutation in experiment E15 is the same as for the single gene mutation in pkaC (experiment E9).

ID	Gene 1	Gene 2	aggregation	Confidence	Comments
E2	yakA-		-	0.50	
E9	pkaC+		++	0.50	
E15	yakA-	pkaC+	++	0.20	

Starting With GenePath

Inference of Genetic Networks

GenePath uses relations inferred from experimental data and relations given in prior knowledge to propose a single genetic network. Many different genetic networks may be consistent with your data and relations that stem from them. But GenePath reports only on the network where it can prove each of its elements, that is, where each relation in the network is either given as the prior knowledge or it can be inferred from the data.

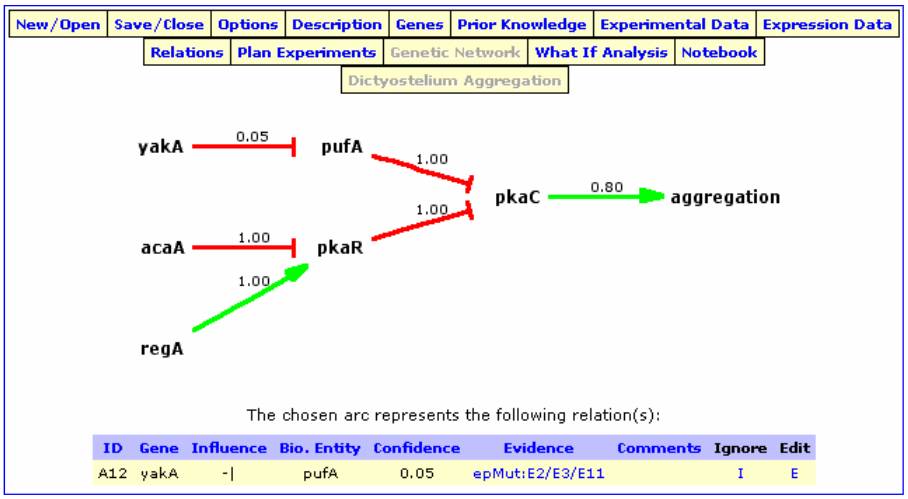

The genetic network inferred from the *Dictyostelium* aggregation data is shown on the right. Edges that represent excitation are shown in green, and those for inhibition are shown in red. For instance, a gene called yakA inhibits pufA, and regA excites pkaR. Numbers that accompany each edge give the confidence for the relation, which was computed from the confidences of experiments that participated in its inference.

Elements of the genetic network — genes and edges that relate them — are clickable: clicking on a gene would show all the data and relations where this gene is involved. Selecting an edge by clicking on its confidence shows the corresponding relations.

For example, the picture of genetic network above is the snapshot of the screen just after we clicked on the edge between yakA and pufA.

Circular Relations, Conflicts

Genetic networks proposed by GenePath are not necessary directed acyclic graphs. Depending on you data, circular relations may be inferred and these are shown in GenePath's networks as cycles. For instance, in the Dictyostelium Communication Adhesion project, the network shows a circular relation between genes lagC and lagD. These two genes are shown in the box, as it could not be determined which one is the one that gene comC has a direct influence on.

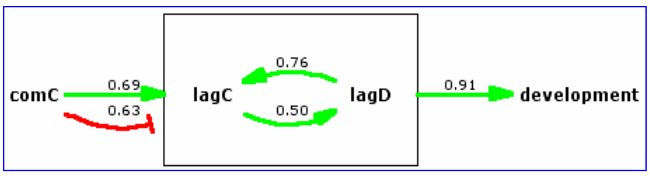

Notice also that for this network, GenePath determined that comC can both either excite or inhibit either of the lag genes. While we may say that these two relations are in conflict, it is often the case that a logical explanation for it may be found and hence GenePath displays both in the network.

GenePath

What-If Analysis

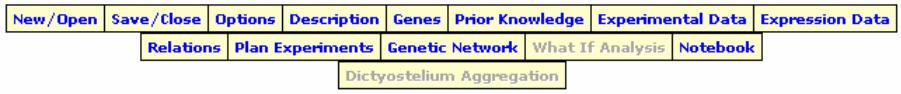

What-if analysis is a simple yet powerful way to explore how each entry in our experimental data set influences the corresponding genetic network. For this, GenePath provides a screen with a genetic network and tables with experimental data and prior knowledge. Ignoring any entry, or changing it by editing would have an immediate affect on the network.

You can use the what-if analysis to determine how robust is your network with respect to a particular experimental finding. You may, for instance, find out that ignoring the experiment does not change the network or, adversely, completely alters its topology. Notice, for instance, how the genetic network for the aggregation of *Dictyostelium* changes after ignoring the last two relations in the prior knowledge!

"You can use the what-if analysis to determine how robust is your network with respect to a particular experimental finding."

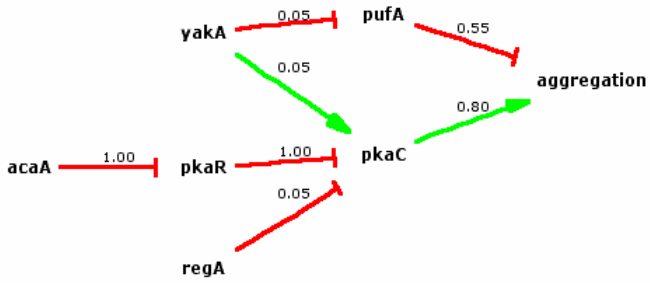

Genetic experiments:

ID	Gene 1	Gene 2	aggregation	Confidence	Comments	Ignore	Edit
E1			+	1.00		I	E
E2	yakA-		-	0.50		I	E
E3	pufA-		++	0.50		I	E
E4	pkaR-		++	0.50		I	E
E5	pkaC-		-	0.50		I	E
E6	acaA-		-	0.50		I	E
E7	regA-		++	0.50		I	E
E8	acaA+		++	0.50		I	E
E9	pkaC+		++	0.50		I	E
E10	pkaC-	regA-	-	0.20		I	E
E11	yakA-	pufA-	++	0.20		I	E
E12	yakA-	pkaR-	±	0.20		I	E
E13	yakA-	pkaC-	-	0.20		I	E
E14	pkaC-	yakA+	-	0.20		I	E
E15	yakA-	pkaC+	++	0.20		I	E

Prior knowledge:

ID	Gene	Influence	Bio. Entity	Confidence	Comments	Ignore	Edit
P1	pkaR	-|	pkaC	1.00		I	E
P2	acaA	-|	pkaR	1.00		I	E
P3	regA	->	pkaR	1.00		I	E
P4	pufA	-|	pkaC	1.00		I	E

Starting With GenePath

Experiment Planning

GenePath can propose experiments to augment those already in the data base. There are in principle many experiments that can be proposed and that would have an effect on the resulting network. To systematically search for potentially interesting experiments, GenePath starts with a matrix of genes which reports on which relations are yet to be defined.

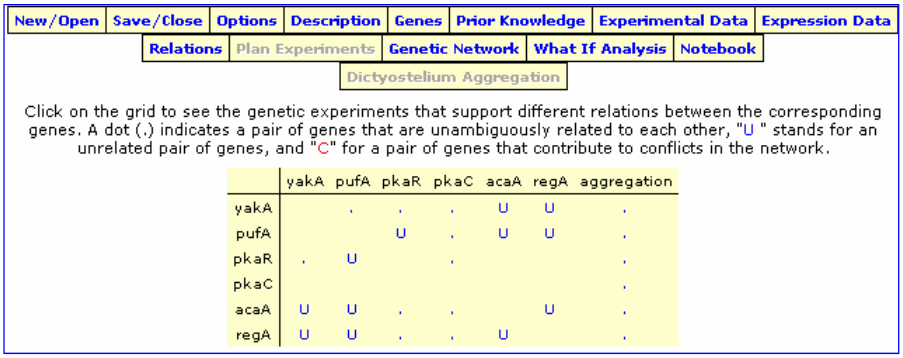

For instance, in the matrix for the *Dictyostelium* aggregation, the relation between regA and acaA (last row, fifth column) is not defined. The possible relations between these two genes (if imposing that regA would be upstream of acaA) include inhibition, excitation, parallelism, and a hypothesis that regA cannot be upstream of acaA. Each such relation may require a set of experiments to be inferred. Out of these experiments, some may be already in your data base, but at least an additional one is missing for the inference. Clicking on the undefined marker — the letter "U" in the matrix — instructs GenePath to search for such experiments and report them accordingly to the type of relations that may benefit from them.

The snapshot below lists the experiments as reported by GenePath that would be needed to hypothesize the inhibitory relation between regA and acaA. There are four experimental plans provided and sorted by increased cost (difficulty) to obtain them. The first two involve only expression experiments: does the expression of acaA change under regA mutant background? Notice that for the first plan, the relation between the two genes would be inhibitory only if the acaA would be overexpressed.

Someehat less obvious is the third plan, which proposes a new experiment involving a double mutant with a knock-out of regA and acaA. This is a proposed experiment, and is distinguished from the two single-mutant experiments already in the data base as it is printed in red. Notice that regA inhibits acaA only if the outcome of this particular experiment is no aggregation ("-").

And finally, there is a fourth plan of experiments, which proposes two additional experiments. Three particular combinations of outcomes — each given in a distinct column under "aggregation" — would lead to a hypothesized relation between regA and acaA. Notice that GenePath uses radio buttons to ask the user which particular combination of outcomes was observed.

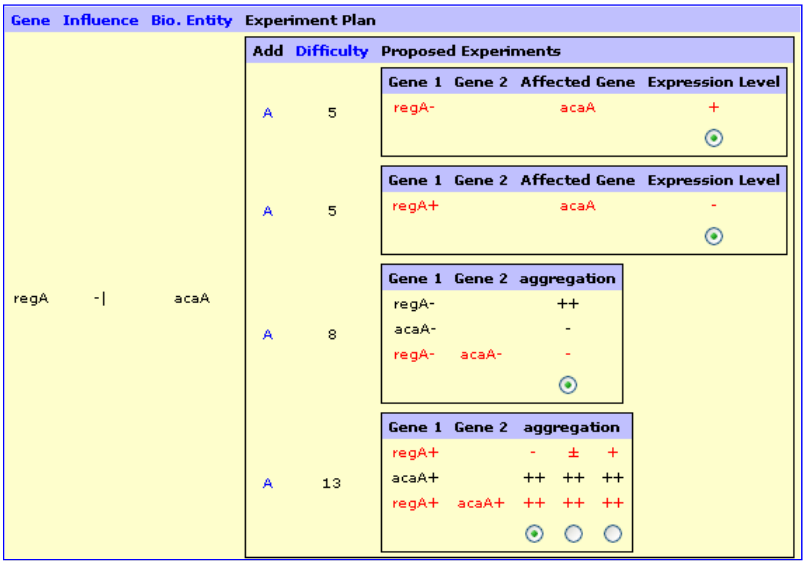

GenePath

Adding Experiments from the Experiment Plan

Experiment planning is not only about reviewing experiments that could be done to infer some interesting relation, but also about examining their effect on the set of previously established relations and pathways. In GenePath, if you agree with any of the experiments proposed by GenePath, you may click on "A" to add these experiments to the project database.

Let us look at an example. For our *Dictyostelium* aggregation project, the relation between pufA and pkaR — with pufA upstream to pkaR — has not been observed in the data (see the matrix on the right). Suppose we agree with a rather expensive experimental plan proposed by GenePath and add the following two experiments to the data: overexpression of pufA with a wild-type aggregation, and overexpression of pufA with knockout of pkaR with excessive aggregation. GenePath reports on the two experiments added, but also on the conflict caused by them in the relation between yakA and pkaR.

	yakA	pufA	pkaR	pkaC	acaA	regA	aggregation
yakA		.	.	.	U	U	.
pufA			U	.	U	U	.
pkaR	.	U		.			.
pkaC							.
acaA	U	U	.	.		U	.
regA	U	U	.	.	U		.

pufA -> pkaR

Add	Difficulty	Proposed Experiments
A	5	Gene 1: pufA- · Gene 2: · Affected Gene: pkaR · Expression Level: -
A	5	Gene 1: pufA+ · Gene 2: · Affected Gene: pkaR · Expression Level: +
A	13	Gene 1: pufA- · Gene 2: · aggregation: ++ Gene 1: pkaR+ · Gene 2: · aggregation: - Gene 1: pufA- · Gene 2: pkaR+ · aggregation: -
A	13	Gene 1: pufA+ · Gene 2: · aggregation: - Gene 1: pkaR- · Gene 2: · aggregation: ++ Gene 1: pufA+ · Gene 2: pkaR- · aggregation: ++

	yakA	pufA	pkaR	pkaC	acaA	regA	aggregation
yakA		.	C	.	U	U	.
pufA			.	.	U	U	.
pkaR	.			.			.
pkaC							.
acaA	U	U	.	.		U	.
regA	U	U	.	.	U		.

The following experiment(s) / expression data has been added:

ID	Gene 1	Gene 2	aggregation	Confidence	Comments
E16	pufA+		+	0.50	hypothetic experiment
E17	pufA+	pkaR-	++	0.20	hypothetic experiment

Abduced relation(s):

ID	Gene	Influence	Bio. Entity	Confidence	Evidence
A15	pufA	->	pkaR	0.05	epMut:E16/E4/E17

Dealing with Conflicts

Conflicts emerge if GenePath infers two or more different relations between the two. Conflicts are reported with a red letter "C" in the matrix used in experiment planning (see above) and can be further investigated by reviewing which are the particular relations that are causing the conflicts. For example, the conflict introduced between yakA and pkaR by two new experiments as added in the example above are due to two different relations: one claims that yakA and pkaR are parallel and infers this directly from the data, the other establishes the inhibitory relation based on the relations that additionally involve pufA. GenePath does not include a mechanism that would somehow automatically resolve such conflicts, but rather provides interfaces to find what caused them and which particular set of experiments is responsible for it. It is then up to the user to either acknowledge and agree with such findings, or resolve the conflict by changing the experimental database.

ID	Gene	Influence	Bio. Entity	Confidence	Evidence
A13	pkaR	||	yakA	0.05	parDiff:E2/E4/E12
A17	yakA	-|	pkaR	0.00	epTC[yakA -| pufA, pufA -> pkaR]

Starting With GenePath

Notebook

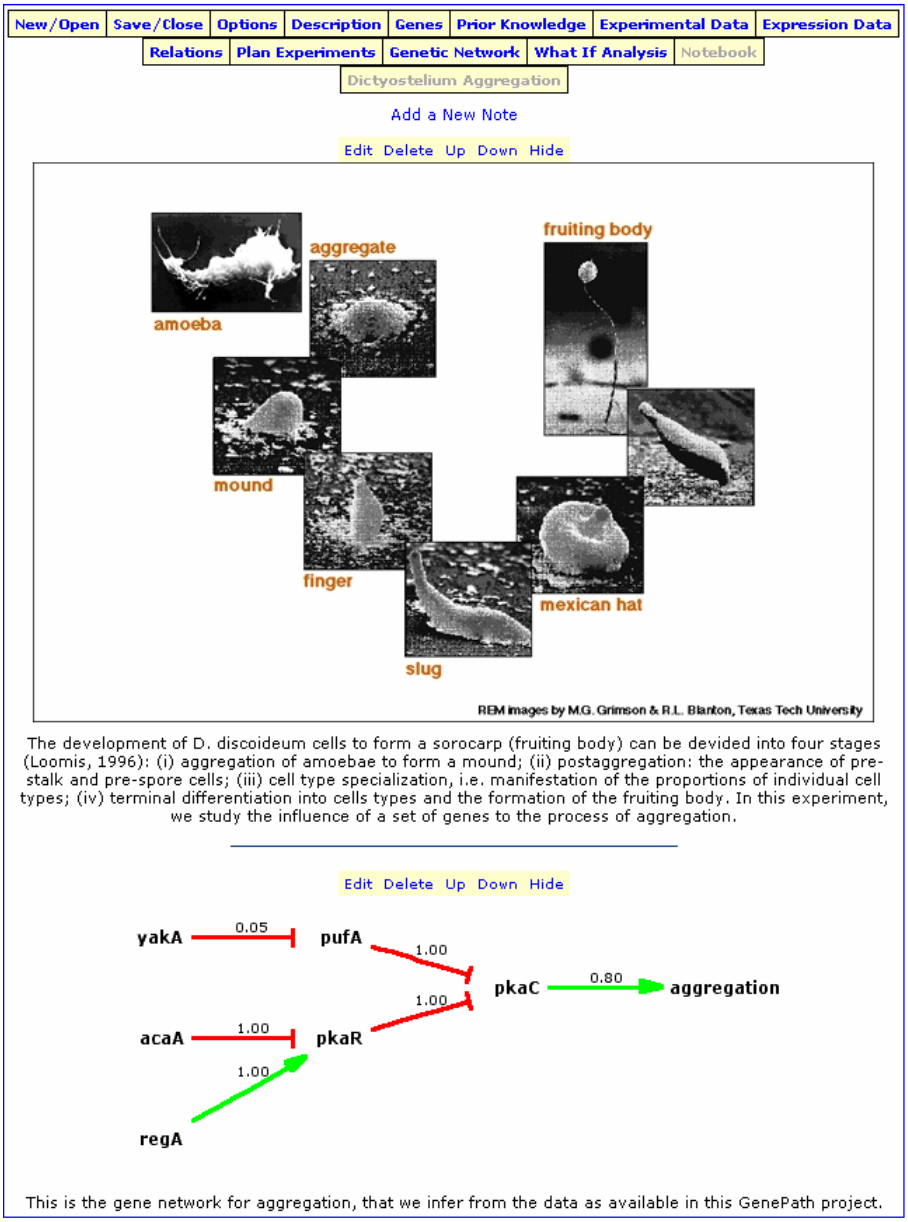

GenePath's notebooks are bitmaps (BMP), GIF and JPEG. To save space and time for communication — GenePath is a server-based system — compressed pictures such as those in JPEG format are preferred.

Also notice that information from the notebook will be save in the same XML file as the rest of the information that composes your project. Notebook text and pictures will be stored at the end of the XML file, maintaining the "comprehensibility" of the rest and (more) important parts of the XML file. But be aware that with any large graphical material the size of the saved project will dramatically increase, thus adding to the time you will need to upload and save your projects to and from GenePath server.

"When using a notebook, to save on memory and time for communication — GenePath is a server-based system — compressed pictures such as those in JPEG format are preferred."

Notebook is a place where you would enter any comments about the particular data, relations, and networks you are using GenePath to analyze. Originally we thought that this is the place were one would place snapshots of networks that evolve when using GenePath for exploratory data analysis, but one can place any graphics within a notebook and provide, for instance, background information on a particular project.

Picture formats supported in

GenePath

Working with the Notebook

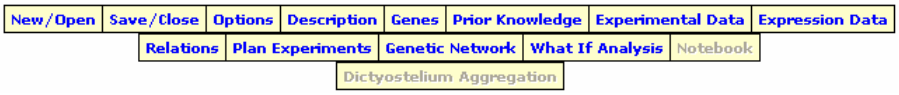

Notebook is composed of entries that include a text and a single picture, which is optional. You may add a new note by pressing "Add a New Note" link in the Notebook window. A dialog will open, and you may either start typing the text (in the entry labeled with "Caption") or adding a picture. The file with a picture should be available somewhere on your local computer: use your browse to specify which file to use and then click append to add the picture to the dialog.

Once finished with writing the comment and selecting the appropriate picture, do not forget to press the "Add" button which adds your new entry at the end of the notebook.

You may change the order of entries in the notebook by using the associated "Up" and "Down" buttons on the top of each entry. Pressing the "Hide" button will hide the picture of the entry and display only the first few words of the text.

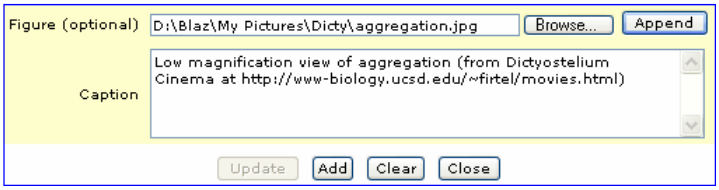

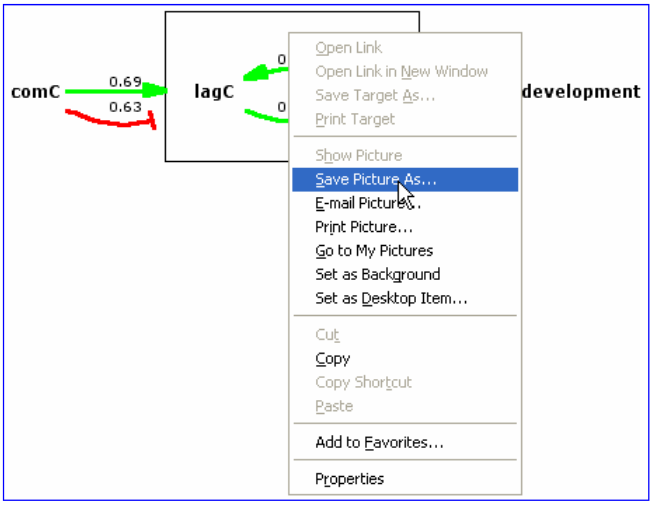

You may often want to add a snapshot of the current network to the notebook. To do this, go to the screen with the network (menu "Genetic Network"), right click anywhere in the network, choose "Save Picture As …" and store the picture somewhere on your computer disk. Now choose the "Notebook", and include the picture you have just saved in the new entry.

Starting With GenePath

Saving Your Work

GenePath is server-based software. This has many advantages, including being able to always work with the most recent version of the software, use it equally on Macs and PCs, or access it from anywhere. But you should also be aware that any faults in the network connections may make you lose your work. Also, the session you will have open at GenePath's server will automatically reset after a grace period of thirty minutes. So, make sure that you save your work frequently, and definitely do so before you go on a coffee break.

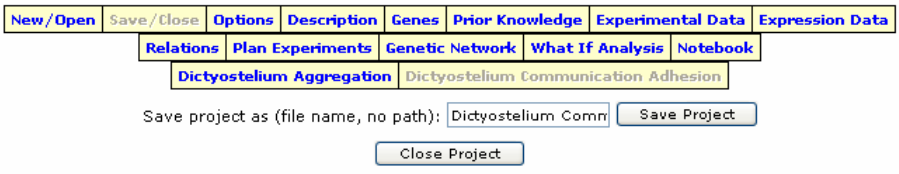

GenePath can encode all the data and information you have embedded in the project within a single XML-based description. To save this in a file, choose the Save/Close button from GenePath's menu and then click on "Save Project". This will prepare the XML description of your data: right click on the link provided, and save it to a desired data file on your local computer.

When you need to upload the data to GenePath, just use the "Browse" button on the project selection screen (see page 2 of this Guide), choose your XML file with the project information, and click "Open".

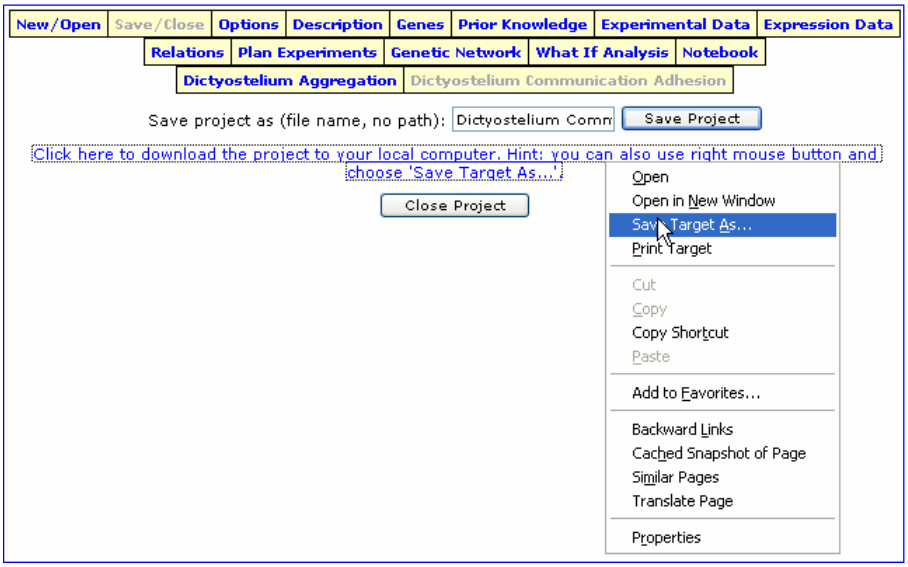

"GenePath maintains all of your data within a session that runs on GenePath's server. To avoid loosing it, make sure you save your work to the local data file frequently."

Baylor College of Medicine & University of Ljubljana

Department of Molecular and Human Genetics
Baylor College of Medicine
One Baylor Plaza
Houston, Texas 77030

Artificial Intelligence Laboratory
Faculty of Computer and Information Science
University of Ljubljana
Trzaska 25
SI-1000 Ljubljana, Slovenia

Further Reading

Zupan B, Bratko I, Demsar J, Juvan P, Halter JA, Kuspa A, Shaulsky G: GenePath: a system for automated construction of genetic networks from mutant data. In Bioinformatics 19(3), pp.383-389, 2003.

Zupan B, Bratko I, Demsar J, Juvan P, Curk T, Borstnik U, Beck JR, Halter J, Kuspa A, Shaulsky G: GenePath: a system for inference of genetic networks and proposal of genetic experiments. In Artificial Intelligence in Medicine 29(1-2), pp.107-130, 2003.

GenePath's Web Site: genepath.org

GenePath Team

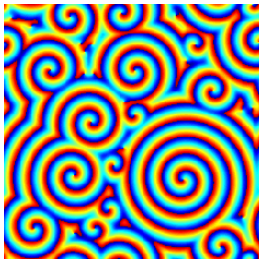

A computer rendering of a spiral wave like the spatial pattern formation during aggregation of the Dictyostelium

GenePath is a result of a collaboration initiated by John Halter between between the Laboratory for Artificial Intelligence at the Faculty of Computer and Information Science, University of Ljubljana, Slovenia, and the Departments of Molecular and Human Genetics and Biochemistry and Molecular Biology at Baylor College of Medicine, Houston, TX.

A number of people that specialize either in biology and genetics or in computer science participate in this project. From the perspective of biology and genetics, GenePath was crafted by Gad Shaulsky and Adam Kuspa. The computer program was first developed by Blaz Zupan and Janez Demsar using a specialized language for crafting artificial intelligence-based applications called Prolog. While the authors liked the elegance of Prolog and its text-based interface, the participating biologists didn't. So Peter Juvan started to work on a web-based interface and added many new features and functions, which lead to the version of GenePath that is reviewed in this Guide.

The particular approach we chose for GenePath was by large influenced by Ivan Bratko, who in addition did preliminary studies in qualitative reasoning and simulations of genetic networks. Experiment planning was first prototyped in Prolog by Tomaz Curk and Urban Borstnik. John Halter and J. Robert Beck provided many useful comments and suggestions. For these, we are also thankful to Chad Shaw, Nancy van Van Driessche, and Victoria Lundblad.

Work on GenePath was supported by the Program and Travel Grant by The Slovene Ministry of Education, Science and Sport, a grant from the National Institute of Child Health and Human Development (P01 HD39691-01), and by a travel grant from the National Academy of Sciences under the Collaboration in Basic Science and Engineering supported by Contract No. INT-0002341 from the National Science Foundation.