# graphtik Documentation

*Release 2.1.1.dev0*

**Yahoo Vision and Machine Learning Team: Huy Nguyen, Arel Cord**

**Nov 12, 2019**

# Contents

2.1.1.dev0, Nov 12, 2019

**It's a DAG all the way down!**

## Lightweight computation graphs for Python

**Graphtik** is an an understandable and lightweight Python module for building and running ordered graphs of computations. The API posits a fair compromise between features and complexity, without precluding any. It can be used as is to build machine learning pipelines for data science projects. It should be extendable to act as the core for a custom ETL engine or a workflow-processor for interdependent files and processes.

**Note:** *Graphtik* is a temporary fork of Graphkit to experiment with Python 3.6+ features.

## 1.1 Operations

At a high level, an operation is a node in a computation graph. Graphtik uses an `operation` class to represent these computations.

### 1.1.1 The `operation` class

The `operation` class specifies an operation in a computation graph, including its input data dependencies as well as the output data it provides. It provides a lightweight wrapper around an arbitrary function to make these specifications.

There are many ways to instantiate an `operation`, and we'll get into more detail on these later. First off, though, here's the specification for the `operation` class:

### 1.1.2 Operations are just functions

At the heart of each `operation` is just a function, any arbitrary function. Indeed, you can instantiate an `operation` with a function and then call it just like the original function, e.g.:

```
>>> from operator import add
>>> from graphtik import operation

>>> add_op = operation(name='add_op', needs=['a', 'b'], provides=['a_plus_b'])(add)

>>> add_op(3, 4) == add(3, 4)
True
```

### 1.1.3 Specifying graph structure: `provides` and `needs`

Of course, each `operation` is more than just a function. It is a node in a computation graph, depending on other nodes in the graph for input data and supplying output data that may be used by other nodes in the graph (or as a graph output). This graph structure is specified via the `provides` and `needs` arguments to the `operation` constructor. Specifically:

- `provides`: this argument names the outputs (i.e. the returned values) of a given `operation`. If multiple outputs are specified by `provides`, then the return value of the function comprising the `operation` must return an iterable.

- `needs`: this argument names data that is needed as input by a given `operation`. Each piece of data named in needs may either be provided by another `operation` in the same graph (i.e. specified in the `provides` argument of that `operation`), or it may be specified as a named input to a graph computation (more on graph computations *here*).

When many operations are composed into a computation graph (see *Graph Composition* for more on that), Graphtik matches up the values in their `needs` and `provides` to form the edges of that graph.

Let's look again at the operations from the script in *Quick start*, for example:

```
>>> from operator import mul, sub
>>> from functools import partial
>>> from graphtik import compose, operation

>>> # Computes |a|^p.
>>> def abspow(a, p):
...     c = abs(a) ** p
...     return c

>>> # Compose the mul, sub, and abspow operations into a computation graph.
>>> graphop = compose(name="graphop")(
...     operation(name="mul1", needs=["a", "b"], provides=["ab"])(mul),
...     operation(name="sub1", needs=["a", "ab"], provides=["a_minus_ab"])(sub),
...     operation(name="abspow1", needs=["a_minus_ab"], provides=["abs_a_minus_ab_cubed
↪"])
...     (partial(abspow, p=3))
... )
```

---

**Tip:** Notice the use of `functools.partial()` to set parameter `p` to a contant value.

---

The `needs` and `provides` arguments to the operations in this script define a computation graph that looks like this (where the oval are *operations*, squares/houses are *data*):

---

**Tip:** See *Plotting* on how to make diagrams like this.

---

### 1.1.4 Instantiating operations

There are several ways to instantiate an `operation`, each of which might be more suitable for different scenarios.

#### Decorator specification

If you are defining your computation graph and the functions that comprise it all in the same script, the decorator specification of `operation` instances might be particularly useful, as it allows you to assign computation graph structure to functions as they are defined. Here's an example:

```
>>> from graphtik import operation, compose

>>> @operation(name='foo_op', needs=['a', 'b', 'c'], provides='foo')
... def foo(a, b, c):
...     return c * (a + b)

>>> graphop = compose(name='foo_graph')(foo)
```

#### Functional specification

If the functions underlying your computation graph operations are defined elsewhere than the script in which your graph itself is defined (e.g. they are defined in another module, or they are system functions), you can use the functional specification of `operation` instances:

```
>>> from operator import add, mul
>>> from graphtik import operation, compose

>>> add_op = operation(name='add_op', needs=['a', 'b'], provides='sum')(add)
>>> mul_op = operation(name='mul_op', needs=['c', 'sum'], provides='product')(mul)

>>> graphop = compose(name='add_mul_graph')(add_op, mul_op)
```

The functional specification is also useful if you want to create multiple `operation` instances from the same function, perhaps with different parameter values, e.g.:

```
>>> from functools import partial

>>> def mypow(a, p=2):
...     return a ** p

>>> pow_op1 = operation(name='pow_op1', needs=['a'], provides='a_squared')(mypow)
>>> pow_op2 = operation(name='pow_op2', needs=['a'], provides='a_cubed
↪')(partial(mypow, p=3))

>>> graphop = compose(name='two_pows_graph')(pow_op1, pow_op2)
```

A slightly different approach can be used here to accomplish the same effect by creating an operation "factory":

```
>>> def mypow(a, p=2):
...     return a ** p

>>> pow_op_factory = operation(mypow, needs=['a'], provides='a_squared')

>>> pow_op1 = pow_op_factory(name='pow_op1')
>>> pow_op2 = pow_op_factory(name='pow_op2', provides='a_cubed', fn=partial(mypow,
↪p=3))

>>> graphop = compose(name='two_pows_graph')(pow_op1, pow_op2)
>>> graphop({'a': 2})
{'a': 2, 'a_cubed': 8, 'a_squared': 4}
```

**Note:** You cannot call again the factory to overwrite the *function*, you have to use the `fn=` keyword.

### 1.1.5 Modifiers on `operation` inputs and outputs

Certain modifiers are available to apply to input or output values in `needs` and `provides`, for example to designate an optional input. These modifiers are available in the `graphtik.modifiers` module:

**Optionals**

**Sideffects**

## 1.2 Graph Composition

Graphtik's `compose` class handles the work of tying together `operation` instances into a runnable computation graph.

### 1.2.1 The `compose` class

For now, here's the specification of `compose`. We'll get into how to use it in a second.

### 1.2.2 Simple composition of operations

The simplest use case for `compose` is assembling a collection of individual operations into a runnable computation graph. The example script from *Quick start* illustrates this well:

```
>>> from operator import mul, sub
>>> from functools import partial
>>> from graphtik import compose, operation

>>> # Computes |a|^p.
>>> def abspow(a, p):
...     c = abs(a) ** p
...     return c

>>> # Compose the mul, sub, and abspow operations into a computation graph.
```

(continues on next page)

```
>>> graphop = compose(name="graphop")(
...     operation(name="mul1", needs=["a", "b"], provides=["ab"])(mul),
...     operation(name="sub1", needs=["a", "ab"], provides=["a_minus_ab"])(sub),
...     operation(name="abspow1", needs=["a_minus_ab"], provides=["abs_a_minus_ab_cubed
↪"])
...     (partial(abspow, p=3))
... )
```

The call here to `compose()` yields a runnable computation graph that looks like this (where the circles are operations, squares are data, and octagons are parameters):

## 1.2.3 Running a computation graph

The graph composed in the example above in *Simple composition of operations* can be run by simply calling it with a dictionary argument whose keys correspond to the names of inputs to the graph and whose values are the corresponding input values. For example, if `graph` is as defined above, we can run it like this:

```
# Run the graph and request all of the outputs.
>>> out = graphop({'a': 2, 'b': 5})
>>> out
{'a': 2, 'b': 5, 'ab': 10, 'a_minus_ab': -8, 'abs_a_minus_ab_cubed': 512}
```

### Producing a subset of outputs

By default, calling a graph-operation on a set of inputs will yield all of that graph's outputs. You can use the `outputs` parameter to request only a subset. For example, if `graphop` is as above:

```
# Run the graph-operation and request a subset of the outputs.
>>> out = graphop({'a': 2, 'b': 5}, outputs="a_minus_ab")
>>> out
{'a_minus_ab': -8}
```

When using `outputs` to request only a subset of a graph's outputs, Graphtik executes only the `operation` nodes in the graph that are on a path from the inputs to the requested outputs. For example, the `abspow1` operation will not be executed here.

### Short-circuiting a graph computation

You can short-circuit a graph computation, making certain inputs unnecessary, by providing a value in the graph that is further downstream in the graph than those inputs. For example, in the graph-operation we've been working with, you could provide the value of `a_minus_ab` to make the inputs `a` and `b` unnecessary:

```
# Run the graph-operation and request a subset of the outputs.
>>> out = graphop({'a_minus_ab': -8})
>>> out
{'a_minus_ab': -8, 'abs_a_minus_ab_cubed': 512}
```

When you do this, any `operation` nodes that are not on a path from the downstream input to the requested outputs (i.e. predecessors of the downstream input) are not computed. For example, the `mul1` and `sub1` operations are not executed here.

This can be useful if you have a graph-operation that accepts alternative forms of the same input. For example, if your graph-operation requires a `PIL.Image` as input, you could allow your graph to be run in an API server by adding an earlier `operation` that accepts as input a string of raw image data and converts that data into the needed `PIL.Image`. Then, you can either provide the raw image data string as input, or you can provide the `PIL.Image` if you have it and skip providing the image data string.

### 1.2.4 Adding on to an existing computation graph

Sometimes you will have an existing computation graph to which you want to add operations. This is simple, since `compose` can compose whole graphs along with individual `operation` instances. For example, if we have `graph` as above, we can add another operation to it to create a new graph:

```
>>> # Add another subtraction operation to the graph.
>>> bigger_graph = compose(name="bigger_graph")(
...     graphop,
...     operation(name="sub2", needs=["a_minus_ab", "c"], provides="a_minus_ab_minus_c
→")(sub)
... )

>>> # Run the graph and print the output.
>>> sol = bigger_graph({'a': 2, 'b': 5, 'c': 5}, outputs=["a_minus_ab_minus_c"])
>>> sol
{'a_minus_ab_minus_c': -13}
```

This yields a graph which looks like this (see *Plotting*):

```
>>> bigger_graph.plot('bigger_example_graph.svg', solution=sol)
```

### 1.2.5 More complicated composition: merging computation graphs

Sometimes you will have two computation graphs—perhaps ones that share operations—you want to combine into one. In the simple case, where the graphs don't share operations or where you don't care whether a duplicated operation is run multiple (redundant) times, you can just do something like this:

```
combined_graph = compose(name="combined_graph")(graph1, graph2)
```

However, if you want to combine graphs that share operations and don't want to pay the price of running redundant computations, you can set the `merge` parameter of `compose()` to `True`. This will consolidate redundant `operation` nodes (based on `name`) into a single node. For example, let's say we have `graphop`, as in the examples above, along with this graph:

```
>>> # This graph shares the "mul1" operation with graph.
>>> another_graph = compose(name="another_graph")(
...     operation(name="mul1", needs=["a", "b"], provides=["ab"])(mul),
...     operation(name="mul2", needs=["c", "ab"], provides=["cab"])(mul)
... )
```

We can merge `graphop` and `another_graph` like so, avoiding a redundant `mul1` operation:

```
>>> merged_graph = compose(name="merged_graph", merge=True)(graphop, another_graph)
>>> print(merged_graph)
NetworkOperation(name='merged_graph',
```

(continues on next page)

```
                    needs=[optional('a'), optional('b'), optional('c')],
                    provides=['ab', 'a_minus_ab', 'abs_a_minus_ab_cubed', 'cab'])
```

This `merged_graph` will look like this:

As always, we can run computations with this graph by simply calling it:

```
>>> merged_graph({'a': 2, 'b': 5, 'c': 5}, outputs=["cab"])
{'cab': 50}
```

# 1.3 Plotting and Debugging

## 1.3.1 Plotting

For *Errors & debugging* it is necessary to visualize the graph-operation. You may plot the original plot and annotate on top the *execution plan* and solution of the last computation, calling methods with arguments like this:

```
graphop.plot(show=True)                 # open a matplotlib window
graphop.plot("graphop.svg")              # other supported formats: png, jpg, pdf, ...
graphop.plot()                          # without arguments return a pydot.DOT object
graphop.plot(solution=out)              # annotate graph with solution values
```

Fig. 1: The legend for all graphtik diagrams, generated by `graphtik.plot.legend()`.

The same `plot()` methods are defined on a `NetworkOperation`, `Network` & `ExecutionPlan`, each one capable to produce diagrams with increasing complexity. Whenever possible, the top-level `plot()` methods delegates to the ones below.

For instance, when a net-operation has just been composed, plotting it will come out bare bone, with just the 2 types of nodes (data & operations), their dependencies, and the sequence of the execution-plan.

But as soon as you run it, the net plot calls will print more of the internals. Internally it delegates to `ExecutionPlan.plot()` of ``graph_op.last_plan`() attribute, which *caches* the last run to facilitate debugging. If you want the bare-bone diagram, plot network:

```
netop.net.plot(...)
```

**Note:** For plots, Graphviz program must be in your PATH, and `pydot` & `matplotlib` python packages installed. You may install both when installing `graphtik` with its `plot` extras:

```
pip install graphtik[plot]
```

**Tip:** The pydot.Dot instances returned by `plot()` are rendered directly in *Jupyter/IPython* notebooks as SVG images.

## 1.3.2 Errors & debugging

Graphs may become arbitrary deep. Launching a debugger-session to inspect deeply nested stacks is notoriously hard

As a workaround, when some operation fails, the original exception gets annotated with the folllowing properties, as a debug aid:

```
>>> from graphtik import compose, operation
>>> from pprint import pprint
```

```
>>> def scream(*args):
...     raise ValueError("Wrong!")
```

```
>>> try:
...     compose("errgraph")(
...         operation(name="screamer", needs=['a'], provides=["foo"])(scream)
...     )({'a': None})
... except ValueError as ex:
...     pprint(ex.jetsam)
{'args': {'args': [None], 'kwargs': {}},
 'executed': set(),
 'network': Network(
    +--a
    +--FunctionalOperation(name='screamer', needs=['a'], provides=['foo'], fn='scream
↪')
    +--foo),
 'operation': FunctionalOperation(name='screamer', needs=['a'], provides=['foo'], fn=
↪'scream'),
 'outputs': None,
 'plan': ExecutionPlan(inputs=('a',), outputs=(), steps:
  +--FunctionalOperation(name='screamer', needs=['a'], provides=['foo'], fn='scream
↪')),
 'provides': ['foo'],
 'results': None,
 'solution': {'a': None}}
```

In interactive *REPL* console you may use this to get the last raised exception:

```
import sys

sys.last_value.jetsam
```

The following annotated attributes *might* have meaningfull value on an exception:

**network** the innermost network owning the failed operation/function

**plan** the innermost plan that executing when a operation crashed

**operation** the innermost operation that failed

**args** either the input arguments list fed into the function, or a dict with both `args` & `kwargs` keys in it.

**outputs** the names of the outputs the function was expected to return

**provides** the names eventually the graph needed from the operation; a subset of the above, and not always what has been declared in the operation.

**results** the values dict, if any; it maybe a *zip* of the provides with the actual returned values of the function, ot the raw results.

**executed`** a set with the operation nodes & instructions executed till the error happened.

Ofcourse you may use many of the above "jetsam" values when plotting.

---

**Note:** The *Plotting* capabilities, along with the above annotation of exceptions with the internal state of plan/operation often renders a debugger session unnecessary. But since the state of the annotated values might be incomple, you may not always avoid one.

---

### 1.3.3 Execution internals

## 1.4 API Reference

### 1.4.1 Package: *graphtik*

### 1.4.2 Module: *base*

### 1.4.3 Module: *operations*

### 1.4.4 Module: *network*

### 1.4.5 Module: *plot*

## 1.5 Graphtik Changelog

### 1.5.1 TODO

See also #1.

- [ ] use a "start-node" to insert input-values in solution
- [ ] typo(test): overriden–>overriDDen
- [ ] support functions with `*args` and `**kwargs`.

### 1.5.2 v2.1.0 (20 Oct 2019): DROP BW-compatible, Restruct modules/API, Plan perfect evictions

The first non pre-release for 2.x train.

- BRAKE API: DROP Operation's `params` - use funtools.partial() instead.
- BRAKE API: DROP Backward-Compatible `Data` & `Operation` classes,
- BRAKE: DROP Pickle workarounds - expected to use `dill` instead.
- break(jetsam): drop "graphtik_' prefix from annotated attribute
- ENH(op): now `operation()` supported the "builder pattern" with `operation.withset()`.
- REFACT: renamed internal package *functional –> nodes* and moved classes around, to break cycles easier, (`base` works as suppposed to), not to import early everything, but to fail plot early if `pydot` dependency missing.
- REFACT: move PLAN and `compute()` up, from `Network --> NetworkOperation`.

---

- ENH(NET): new PLAN BULDING algorithm produces PERFECT EVICTIONS, that is, it gradually eliminates from the solution all non-asked outputs.

    - enh: pruning now cleans isolated data.

    - enh: eviction-instructions are inserted due to two different conditions: once for unneeded data in the past, and another for unused produced data (those not belonging typo the pruned dag).

    - enh: discard immediately irrelevant inputs.

- ENH(net): changed results, now unrelated inputs are not included in solution.

- refact(sideffect): store them as node-attributes in DAG, fix their combination with pinning & eviction.

- fix(parallel): eviction was not working due to a typo 65 commits back!

### 1.5.3  v2.0.0b1 (15 Oct 2019): Rebranded as *Graphtik* for Python 3.6+

Continuation of #30 as #31, containing review-fixes in huyng/graphkit#1.

#### Network

- FIX: multithreaded operations were failing due to shared `ExecutionPlan.executed`.

- FIX: prunning sometimes were inserting plan string in DAG. (not `_DataNode`).

- ENH: heavily reinforced exception annotations ("jetsam"):

    - FIX: (8f3ec3a) outer graphs/ops do not override the inner cause.

    - ENH: retrofitted exception-annotations as a single dictionary, to print it in one shot (8f3ec3a & 8d0de1f)

    - enh: more data in a dictionary

    - TCs: Add thorough TCs (8f3ec3a & b8063e5).

- REFACT: rename *Delete*–>'Evict', removed *Placeholder* from nadanodes, privatize node-classes.

- ENH: collect "jetsam" on errors and annotate exceptions with them.

- ENH(sideffects): make them always DIFFERENT from regular DATA, to allow to co-exist.

- fix(sideffects): typo in add_op() were mixing needs/provides.

- enh: accept a single string as *outputs* when running graphs.

#### Testing & other code:

- TCs: *pytest* now checks sphinx-site builds without any warnings.

- Established chores with build services:

    - Travis (and auto-deploy to PyPi),

    - codecov

    - ReadTheDocs

### 1.5.4 v1.3.0 (Oct 2019): New DAG solver, better plotting & "sideffect"

Kept external API (hopefully) the same, but revamped pruning algorithm and refactored network compute/compile structure, so results may change; significantly enhanced plotting. The only new feature actually is the `sideffect`` modifier.

#### Network:

- FIX(#18, #26, #29, #17, #20): Revamped DAG SOLVER to fix bad pruning described in #24 & #25

  Pruning now works by breaking incoming provide-links to any given intermedediate inputs dropping operations with partial inputs or without outputs.

  The end result is that operations in the graph that do not have all inputs satisfied, they are skipped (in v1.2.4 they crashed).

  Also started annotating edges with optional/sideffects, to make proper use of the underlying `networkx` graph.


- REFACT(#21, #29): Refactored Network and introduced `ExecutionPlan` to keep compilation results (the old `steps` list, plus input/output names).

  Moved also the check for when to evict a value, from running the execution-plan, to whenbuilding it; thus, execute methods don't need outputs anymore.

- ENH(#26): "Pin* input values that may be overriten by calculated ones.

  This required the introduction of the new `_PinInstruction` in the execution plan.

- FIX(#23, #22-2.4.3): Keep consistent order of `networkx.DiGraph` and *sets*, to generate deterministic solutions.

  *Unfortunately*, it non-determinism has not been fixed in < PY3.5, just reduced the frequency of spurious failures, caused by unstable dicts, and the use of subgraphs.

- enh: Mark outputs produced by `NetworkOperation`'s needs as `optional`. TODO: subgraph network-operations would not be fully functional until *"optional outpus"* are dealt with (see #22-2.5).

- enh: Annotate operation exceptions with `ExecutionPlan` to aid debug sessions,

- drop: methods `list_layers()`/`show layers()` not needed, `repr()` is a better replacement.

#### Plotting:

- ENH(#13, #26, #29): Now network remembers last plan and uses that to overlay graphs with the internals of the planing and execution:

  - execution-steps & order

  - evict & pin instructions

  - given inputs & asked outputs

  - solution values (just if they are present)

  - "optional" needs & broken links during pruning

- REFACT: Move all API doc on plotting in a single module, splitted in 2 phases, build DOT & render DOT

- FIX(#13): bring plot writing into files up-to-date from PY2; do not create plot-file if given file-extension is not supported.

- FEAT: path pydot library to support rendering in *Jupyter notebooks*.

**Testing & other code:**

- Increased coverage from 77% –> 90%.

- ENH(#28): use `pytest`, to facilitate TCs parametrization.

- ENH(#30): Doctest all code; enabled many assertions that were just print-outs in v1.2.4.

- FIX: `operation.__repr__()` was crashing when not all arguments had been set - a condition frequtnly
  met during debugging session or failed TCs (inspired by @syamajala's 309338340).

- enh: Sped up parallel/multihtread TCs by reducing delays & repetitions.

---

**Tip:** You need `pytest -m slow` to run those slow tests.

---

**Chore & Docs:**

- FEAT: add changelog in `CHANGES.rst` file, containing flowcharts to compare versions `v1.2.4 <--> v1.3..0`.

- enh: updated site & documentation for all new features, comparing with v1.2.4.

- enh(#30): added "API reference' chapter.

- drop(build): `sphinx_rtd_theme` library is the default theme for Sphinx now.

- enh(build): Add `test` *pip extras*.

- sound: https://www.youtube.com/watch?v=-527VazA4IQ, https://www.youtube.com/watch?v=8J182LRi8sU&t=43s

## 1.5.5 v1.2.4 (Mar 7, 2018)

- Issues in pruning algorithm: #24, #25

- Blocking bug in plotting code for Python-3.x.

- Test-cases without assertions (just prints).

## 1.5.6 1.2.2 (Mar 7, 2018, @huyng): Fixed versioning

Versioning now is manually specified to avoid bug where the version was not being correctly reflected on pip install
deployments

## 1.5.7 1.2.1 (Feb 23, 2018, @huyng): Fixed multi-threading bug and faster compute through caching of *find_necessary_steps*

We've introduced a cache to avoid computing find_necessary_steps multiple times during each inference call.

This has 2 benefits:

- It reduces computation time of the compute call
- It avoids a subtle multi-threading bug in networkx when accessing the graph from a high number of threads.

### 1.5.8 1.2.0 (Feb 13, 2018, @huyng)

Added *set_execution_method('parallel')* for execution of graphs in parallel.

### 1.5.9 1.1.0 (Nov 9, 2017, @huyng)

Update setup.py

### 1.5.10 1.0.4 (Nov 3, 2017, @huyng): Networkx 2.0 compatibility

Minor Bug Fixes:

- Compatibility fix for networkx 2.0
- *net.times* now only stores timing info from the most recent run

### 1.5.11 1.0.3 (Jan 31, 2017, @huyng): Make plotting dependencies optional

- Merge pull request #6 from yahoo/plot-optional
- make plotting dependencies optional

### 1.5.12 1.0.2 (Sep 29, 2016, @pumpikano): Merge pull request #5 from yahoo/remove-packaging-dep

- Remove 'packaging' as dependency

### 1.5.13 1.0.1 (Aug 24, 2016)

### 1.5.14 1.0 (Aug 2, 2016, @robwhess)

First public release in PyPi & GitHub.

- Merge pull request #3 from robwhess/travis-build
- Travis build

# CHAPTER 2

# Quick start

Here's how to install:

```
pip install graphtik
```

OR with dependencies for plotting support (and you need to install Graphviz program separately with your OS tools):

```
pip install graphtik[plot]
```

Here's a Python script with an example Graphtik computation graph that produces multiple outputs (a * b, a - a * b, and abs(a - a * b) ** 3):

```python
>>> from operator import mul, sub
>>> from functools import partial
>>> from graphtik import compose, operation

# Computes |a|^p.
>>> def abspow(a, p):
...     c = abs(a) ** p
...     return c
```

Compose the mul, sub, and abspow functions into a computation graph:

```python
>>> graphop = compose(name="graphop")(
...     operation(name="mul1", needs=["a", "b"], provides=["ab"])(mul),
...     operation(name="sub1", needs=["a", "ab"], provides=["a_minus_ab"])(sub),
...     operation(name="abspow1", needs=["a_minus_ab"], provides=["abs_a_minus_ab_cubed
→"])
...     (partial(abspow, p=3))
... )
```

Run the graph-operation and request all of the outputs:

```python
>>> graphop({'a': 2, 'b': 5})
{'a': 2, 'b': 5, 'ab': 10, 'a_minus_ab': -8, 'abs_a_minus_ab_cubed': 512}
```

Run the graph-operation and request a subset of the outputs:

```
>>> graphop({'a': 2, 'b': 5}, outputs=["a_minus_ab"])
{'a_minus_ab': -8}
```

As you can see, any function can be used as an operation in Graphtik, even ones imported from system modules!